

Reasoning with Legal Cases: Analogy or Rule Application?

Katie Atkinson

Department of Computer Science, University of Liverpool
Liverpool, UK
K.M.Atkinson@liverpool.ac.uk

Trevor Bench-Capon

Department of Computer Science, University of Liverpool
Liverpool, UK
tbc@csc.liv.ac.uk

ABSTRACT

Modelling reasoning with precedents has been a central concern of AI and Law since its inception. A recent paper has provided a discussion (in jurisprudential terms) of whether such reasoning is best seen as rule application or analogy. We review some of the prominent AI and Law approaches and find that over the years there has been a move away from analogy to rule application. Even in those approaches which do use analogy, however, the analogies handled concern only analogies between cases represented as sets of factors, and do not consider analogies between the elements of the fact situations peculiar to particular cases. In actual practice, however, analogies are used to determine which factors are relevant in a case, and which party is favoured by particular aspects of the case situation. Such analogies relate not to factors, but to real-world elements of the case and are hard to make and critique without a comprehensive common sense ontology. Thus while we may be able to construct specific ontologies to model past examples of such analogical reasoning, which can be useful for simulation and teaching, the ability to perform analogical reasoning on novel situations is, and is likely to remain, infeasible. This conclusion suggests that there will always be limits to our ability to construct systems able to handle new cases presenting novel situations.

CCS CONCEPTS

• Applied computing → Law.

KEYWORDS

argumentation, analogy, reasoning with cases

ACM Reference Format:

Katie Atkinson and Trevor Bench-Capon. 2019. Reasoning with Legal Cases: Analogy or Rule Application?. In *Seventeenth International Conference on Artificial Intelligence and Law (ICAIL '19)*, June 17–21, 2019, Montreal, QC, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3322640.3326695>

1 INTRODUCTION

Artificial Intelligence has been pursued with a variety of different motivations. Some researchers pursue “strong” AI in which the goal

is to build a “person or, more humbly, an animal” [21]. Others pursue “weak” AI, in which the goal is to produce behaviour which, if done by a human, would be considered intelligent [41]. Still others are interested in cognitive science, which uses computer models to explore the mind and its processes. In the latter it is *how* things are done, not *what* is done that is important. Similarly AI and Law has been pursued with a variety of motivations. Some have indeed thought that it would become possible to replace judges with machines [25], although this ambition may be more often attributed to AI and Law researchers by the press¹ than actually pursued by the researchers themselves. Others have simply preferred to build applications intended to support lawyers in matters such as conceptual retrieval [50]. Others have aimed to produce computer assisted learning systems [5]. Still others have been interested in modelling the reasoning processes of lawyers [46].

All these motivations can be seen in AI systems concerned with reasoning with legal precedents. We have systems intended to predict the outcomes of new cases (e.g. [20] and [4]); systems to retrieve relevant cases (e.g. [27]); systems for teaching law students (e.g. [5]) and systems for modelling the reasoning of lawyers (e.g. [48]). In pursuing the last of these aims, the question arises as to whether reasoning with precedents is best seen as the construction of a theory and the application of the rules of that theory, or as reasoning by analogy using previously decided cases. This question was the concern of a recent paper by Katharina Stevens [49], written from a jurisprudential standpoint, which we will use as the starting point for this paper.

In this paper we will try to shed light on which of these various aims appear to be feasibly addressed by AI, especially symbolic AI. In particular we shall distinguish between the *ex post* simulation of legal reasoning in particular cases, useful for both modelling the reasoning processes and training law students, and the *ex ante* reasoning required to predict the outcomes of cases in which novel situations, which do not unequivocally match any existing precedents, may occur. The latter include cases where the rules “run out” [51], and so precedents often fail to provide sufficient guidance. Similarly machine learning approaches can only use past cases, and so may be unable to adapt to novel situations. For novel situations, we suggest that the use of analogy becomes essential, and in such cases the analogy is typically required to be not between *cases* themselves, but between *elements* in the factual situations of the cases, and so draws on common-sense knowledge about the world rather than purely legal knowledge about case law.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICAIL '19, June 17–21, 2019, Montreal, QC, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6754-7/19/06... \$15.00

<https://doi.org/10.1145/3322640.3326695>

¹For example the 2016 *Guardian* article *Artificial intelligence 'judge' developed by UCL computer scientists* <https://www.theguardian.com/technology/2016/oct/24/artificial-intelligence-judge-university-college-london-computer-scientists>. This article relates to [4]. The paper itself made rather more modest claims.

2 REASONING WITH PRECEDENTS

In a recent paper [49] Stevens explores the question of whether “the process of reasoning through which a judge determines whether the existence of a specific precedent gives her a binding authoritative reason to follow in her present-case” is best seen as rule application or analogical reasoning. Stevens identifies the key difference between the two as follows: rule application approaches

“describe reasoning by precedent as a process in which a judge *first* determines a rule from the precedent-case alone and *then* applies it to the present-case. She reasons with each case independently and connects the two through the rule.” p8, *Italics* ours.

In contrast, analogical reasoning is a process

“in which a judge reasons with both cases at the same time, finding similarities and differences between them.” [49], p8.

Roth [47] makes a similar distinction. His thesis, however, compares cases as sets of factors, and does not consider analogies leading to the ascription of factors.

The strength of the rule application approach is that it provides a good mechanism to show how precedents constrain decisions. The weakness is that it does not provide an effective way of distinguishing cases: distinguished cases would be covered by a more specific rule, but there is no mechanism to propose exceptions when no more specific rule exists. Analogical reasoning is the reverse since it does consider the differences between the cases and so easily accounts for distinguishing, but offers a less satisfactory account of precedential constraint. In practice the distinction may become blurred: Stevens suggests

“that both approaches are viable in principle because they can be amended to meet the objections. Once amended, both approaches can accommodate the practice of distinguishing and account for the constraining effects of precedents.” p2.

In the light of this distinction between analogy and rule application, we will begin by considering some leading previous approaches to reasoning with cases in AI and Law, to see the extent to which they can be seen as using analogical reasoning.

3 REASONING WITH PRECEDENT IN AI AND LAW

Reasoning with precedents has been a central topic of AI and Law from the beginning. An overview of the development since 1987 (the year of the first ICAIL) can be found in [14]. A highly influential early approach, upon which subsequent approaches draw heavily, was Rissland and Ashley’s HYPO, introduced in [46] and most fully described in [7]. Our consideration of the development of reasoning with precedents in AI and Law will show a movement from analogical reasoning to rule application.

- *HYPO* [7]. This uses analogical reasoning, because the current case plays a central role. For each new case, a *case lattice* is created, rooted in the current case, and with its children comprising cases sharing aspects² with the current case.

²We follow Stevens in using *aspects* to refer to legally relevant features of the cases. These include both the dimensions of HYPO and the factors of CATO and subsequent

Those with the desired outcome may be cited, those with the opposite outcome can be used as counter-examples. HYPO is also concerned with distinguishing cases: distinguishing and counter-examples are how the cited precedents may be challenged.

- *CATO* [5]. This was a development from HYPO and can also be considered to use analogy. It too starts with a case lattice, and places even more emphasis on distinguishing since its main purpose is to assist law students in recognising which differences between the current case and the precedent represent effective distinctions. To this end it organises factors into a hierarchy to allow for the downplaying of poor distinctions and the emphasis of good ones.
- *Partitioning approaches* [11] and [52]. These approaches also begin with the current case and a precedent. They partition the aspects present in these two cases into: plaintiff aspects in common, defendant aspects in common, aspects which make the current case weaker and aspects which make the precedent case stronger. In [52] the partitions are then used to instantiate a set of argument schemes, intended to support CATO-style reasoning. Again there is an attempt to identify the differences between the cases and the legal significance of these differences.
- *Rewriting as sets of rules* [40]. Prakken and Sartor’s paper represented a significant turning point in the treatment of reasoning with cases, as it showed how the knowledge embodied in precedents could be expressed as a set of rules. Each precedent is represented as three rules. Two are defeasible: the first is that there should *prima facie* be a decision for the plaintiff if all the pro-plaintiff aspects in the precedent are present in the new case; the second is that there should *prima facie* be a decision for the defendant if all the pro-defendant aspects in the precedent are present in the new case. The third expresses a priority between these rules, depending on which party won the precedent case. This approach produces a theory, constructed solely on the basis of the precedents, without reference to the current case, and so very much represents rule application.
- *Theory construction* [17]. One problem with the preceding approach is that strictly it can only justify *a fortiori* reasoning, whereas in practice new cases will present new combinations of aspects, for which no priority rules exist. This limits the usefulness of the theory: it cannot resolve conflicts between rules which have not previously been decided. In [17] an attempt to remedy this was made by associating aspects with the promotion of particular social values, and using precedents to establish preferences between values which could then be used to determine as yet untested combinations of aspects, provided they relate to the same values as featured in the precedents. This approach makes no intrinsic use of the current case, but an empirical realisation of the approach [23] constrained the proposed theories by the need to provide the

systems. As such they may be Boolean (like factors) or have magnitude (like dimensions). They may (like factors) or may not (like dimensions) intrinsically favour a particular party to the case. Moreover they may be derived from facts, as in HYPO, or the cases may be represented as sets of aspects, as in CATO and most subsequent approaches.

desired outcome in the current case. Even so, this implementation was interested in similarities rather than differences, and so cannot be considered analogical in Stevens' sense.

- *Formalisations of precedential constraint* [31] and [42]. These papers build on the rule based rewriting of [40]. In [31] the problem of *a fortiori* is addressed by allowing the rule for the winning side to be broadened so that it contains only a subset of the aspects available in the precedent. The claim is that this represents the reason model of precedential constraint [35] as opposed to the results model [6], which was essentially what was used in [40]. Some refinements are presented in [42].
- *Methodology based on Abstract Dialectical Frameworks* [3]. The ANGELIC methodology presented in [3] is very firmly in the rule application camp. It sets out to encapsulate the knowledge enshrined in a set of cases in a legal domain, for application to new cases. Thus the theory is produced before any new case has arisen, and so there is no possibility of analogical reasoning.

Although these approaches can be seen as developments in a common tradition [14], there is a movement from what Stevens would term an analogical approach, tailored to the current case, to a rule based approach. In the earlier systems the current case was prominent and individual precedents and their differences and distinguishing were at centre stage. By the later systems, however, the theory was formed from the precedents before the new case arose and the nature of precedential constraint and the encapsulation of the knowledge represented by the whole set of precedents and of the priorities between them became central.

Stevens herself favours analogy, believing

“that integrating analogical mapping into reasoning by precedent brings with it important advantages that are lost in reasoning according to the pure rule approach. Specifically, analogical mapping has the benefit of guiding the attention of the judge to the most plausible ways in which she can understand the precedent so that it applies to the present-case.” p39.

According to Stevens there are two advantages to analogy:

“Arguments by analogy are used not only because of their justificatory power, but also because of their rhetorical strength. Analogies are rhetorically effective because they encourage those who are trying to understand them to *restructure* their understanding of the analogs in a way that makes them appear more similar. ... The rule based approach only allows difference in terms drawn from a *pre-identified* list of factors, whereas in analogy, *anything* may serve as a relevant difference.” p34, (*italics* ours)

Should we therefore consider reviving some of the earlier ideas, returning the current case to centre stage, especially if we want to be able to deal with cases which are not straightforwardly decidable from the current precedents?

4 COMPARISON OF THE APPROACHES

Stevens argues that both approaches are viable, once they have been modified to meet the basic objections (the inability of analogy

to capture the notion of precedential constraint and the inability of rules to adequately account for distinguishing). She expresses these modified versions in the form of argumentation schemes and possible attacks on such arguments in the form of questions. For the rule based scheme the questions given in [49], p14. are:

- R1 Is the best interpretation of the opinion of precedent-case *A* that it provides rule *a*?
- R2 Is the legal syllogism of the form *modus ponens* that uses rule *a* in the major premise and case *B* in the minor premise deductively valid?
- R3 Is there any reason that defeats the presumption for reaching decision *Y* based on a property of the present-case *B* that the precedent-case *A* did not have?
- R4 Does case *B* have properties *n1*, *n2*? [i.e. is the antecedent of rule *a* satisfied?].

For analogy the questions are ([49], p32):

- A1 Are *A* and *B* similar in a legally relevant way?
- A2 Can a successful mapping be made to an aspect of the present-case for every aspect of the precedent-case that the opinion highlights enough to indicate that it is part of the *ratio* in the opinion?
- A3 Does the surrounding law allow present-case and precedent-case to be mapped successfully?
- A4 Are there no legally relevant differences between *A* and *B*? Note that, unlike R3 and R4 it makes no difference which case has the additional aspect.

Notable (from a computational point of view) is that the questions relating to rule application are very much more precise than those relating to analogy. For rule application (modelled as in [31]):

- R4 simply checks that all the antecedents of the rule are present in the case, so that the rule is applicable.
- R2 checks that the rule has been correctly applied.
- R3 checks that there is no applicable rule (exception) with the opposite outcome, and with a higher priority.
- R1 relates to the fact that the reason model as expressed in [31] uses a subset of the aspects from the precedent, and so raises the possibility that there may be a more suitable subset.

Thus R2, R3 and R4 are simply a matter of whether the rule has been properly applied. R1 does, however, give some scope for a substantial challenge. Since several subsets of the winning aspects in the precedent could have been chosen as the rule, it is possible to argue that a different subset should have been chosen. If this is so, and the present-case may lack the aspects required to satisfy the new rule, the argument fails.

The questions for analogy are less straightforward to conceive in computational terms. A4 may well be what is addressed by emphasising and downplaying distinctions in CATO, with legal relevance being captured by the notion of downplaying distinctions. A1 can be answered affirmatively if there are aspects in common between the present case and the precedent, and the precedent cannot be distinguished or trumped by a counterexample. A2 could be seen as checking that the required aspects were present in both cases. A3, however, seems to go beyond and suggest that there may be factors which require consideration, but which cannot be found in the existing set of precedents. This suggests that analogy may be

essential when we are confronted with a *novel* case, with relevant aspects not found in any of the previous cases. We will explore this further in the following sections.

5 ANALOGY IS REQUIRED IN “HARD CASES”

A difficulty with the rule model, not discussed by Stevens, but recognised very early in AI and Law [51], is that the rules will run out, and the new case may well fall under two conflicting rules, or under no rule at all. Such cases were termed *hard cases* in [51].

Conflicts can possibly be resolved using priorities as in [40] or value preferences as in [17], or prevented from arising through the precedential constraint mechanism of [31]. The possibility of gaps, however, is harder to avoid. If one only allows precedents to apply to cases which are at least as strong (as in [40] or the results model of [31]), gaps will inevitably be frequent. It is important that rules are capable of going beyond *a fortiori* reasoning from the precedents if they are to provide guidance for a reasonable number of cases. This was the role of values in [17]³, and was the motivation for adopting the reasons model by using only a subset of the aspects favouring the winning side in [31]. Even so, cases will arise which are not covered by any of the established value preferences, or any of the rules, even if they are broadened as suggested in [31].

The problem is exacerbated if we consider that some aspects can have magnitude, as discussed in [29], [43], [16] and [30]. If we adopt these approaches, we need to determine, for all such aspects, a “switching point”⁴ at which the aspect ceases to favour the plaintiff and begins to favour the defendant. Although the switching point is constrained by precedents setting upper and lower bounds, a new case will often fall into the gap between them so that it can be argued by both sides that it is favourable to them (or to neither). If there are several aspects of the new case which have gaps, the case becomes even harder. To resolve such cases, rules cannot help us: we must resort to some kind of analogical reasoning. Note that the gap here is not simply because we have no applicable rule given the set of factors, but that we are even not able to determine which factors are present.

The idea is now to collect those precedents relevant to the new cases (cf the “near misses” of [48]). If the new case falls into a gap these will be the closest precedents, that is those supplying the current upper and lower bounds, and one will favour the plaintiff and the other the defendant. Analogy will then be used to determine which of the bounds the case under consideration should be considered more similar to. Where several aspects with magnitude are involved, there may be several such pairs of precedents, each requiring the determination of a switching point, or of a trade off equation as in [16]. For Boolean aspects, the closest precedents can be determined by looking for analogies to the relevant elements in the current case. Note that, as we shall discuss in more detail below, the analogy in such cases is not between sets of factors, but between factual constituents of the cases, which decide whether particular factors are present or absent in the case or, for aspects with magnitude and

dimensions, which party is favoured in the case. We may therefore distinguish different ways in which a case may be “hard”.

- It may be that the aspects present in the case do not allow the application of any rule derivable from the precedents. In such a case a new preference between the sets of plaintiff and defendant aspects must be determined [39]. Here any analogy can be considered as being between cases represented as sets of aspects, in the manner of [7] and [5].
- It may be that we have an aspect with magnitude which in the case occupies a part of the range not yet determined as being pro-plaintiff or pro-defendant. Here we need to determine which side of the switching point the aspect occupies in this case, and this may involve the use of analogy to determine the most similar precedent. Such analogies will be in terms of the relevant fact situations.
- It may be that we are unsure as to whether an aspect is present or absent. In *Dillon v Legg*, discussed in the next section, the aspect is *close relationship* but it is debatable whether there is a close relationship or not. Again we need to consider analogies between the elements present in the fact situations of the current and precedent cases.

Thus only some hard cases turn on case comparisons at the aspect level. The last two of the above bullets turn on analogies between real-world elements of the cases. There the precedents supply points of reference to be used in the analogy, but no guidance as to whether or not the analogy should be accepted: that is left for the judge to decide (and judges are permitted to disagree, which is often the source of minority opinions). Thus encoding as rules in advance of the decisions is not possible in these situations.

6 ANALOGY AND THE ASCRIPTION OF ASPECTS

One feature common to all the above AI approaches is that a case is seen as capable of being represented as a set of aspects. Although HYPO does compute the dimensions on the basis of facts (in a determinate way, so that there is no room for arguments as to which dimensions apply to a given set of facts), the precedential reasoning only begins when the case lattice is formed on the basis of the applicable dimensions. The other approaches unashamedly represent cases as sets of factors. But this avoids a great deal of interpretation. As Stevens says

“The difficulties in interpreting opinions have been discussed often and at length. Duxbury in [26], at ch. 3, gives an overview over all the problems with interpreting opinions and describes in detail the discussion about the determination of the *ratio decidendi* that took place in the 1930s to 1960s. ... In Horty’s, Bench-Capon’s, and Rigoni’s approaches, these difficulties are bracketed by the authors, who *begin* their theorizing only *after* the aspects of the case are categorized and determined as belonging to the rule’s antecedent or the list of defeated reasons.” p16, footnote 46, (*italics ours*)

This criticism, that the crucial decision of which aspects to ascribe to the cases is taken out of consideration by the system, applies

³The idea in [17] was that, since several factors can promote the same value, a value preference established by a precedent could also be applied to cases where that preference was instantiated by different factors relating to those same values.

⁴Following [43], we use “point”, although often there may be an interval in which neither side is favoured by the aspect.

equally to the other AI and Law approaches as well as those specifically mentioned by Stevens. Once we have identified the set of aspects applicable to the cases, the analogy is between the cases, considered only as these sets of aspects. But examination of actual examples in legal opinions suggests that analogy is often intended to *establish* the similarity between two elements of the fact situation, so that the appropriate aspects can be identified.

Steven's main example for discussing analogy is *Dillon v Legg*⁵ which was a case decided by the Supreme Court of California that established the tort of negligent infliction of emotional distress. It was held that a bystander who suffers damages by the conduct of a negligent *tortfeasor* can recover for negligent infliction of emotional distress. The reasoning, expressed in *Tobriner's* majority opinion, was

“Instead of applying a strict rule, courts should consider factors such as whether the plaintiff was near the scene of the accident, how closely the plaintiff was related to the victim, and what caused the emotional distress.”

Thus we cannot apply the rule until we have identified the set of aspects, such as whether the claimant was sufficiently closely related to the victim, that are applicable to the case. Stevens writes:

“Imagine, for example, that I was to claim that *Dillon v. Legg* should be followed if an accident was not witnessed by the mother of the victim, but instead a kindergarten teacher. I might justify this by saying ‘kindergarten teachers are like mothers.’ In this example, ‘kindergarten teachers’ is the target [of the analogy] while ‘mothers’ is the source.” p21.

The issue concerned is whether the plaintiff is sufficiently closely related to the victim, and the analogy is used to determine whether this is so or not, and so whether the aspect is present in the case. This in turn will determine whether the rule should be applied.

At first sight a kindergarten teacher does not seem to be covered by the decision which speaks of the plaintiff being “closely related” to the victim, since we might understand *relation* as *blood relation* and so see the kindergarten teacher as not related at all. But Stevens offers grounds for an analogy on the basis that both mothers and kindergarten teachers love their charges *because* they have taken care of them. Quite possibly the analogy will be rejected because kindergarten teachers tend to have more wards, and because the care provided is less intimate, so that there are significant differences sufficient to defeat the analogy. But the analogy may be more persuasive in the case of a wet nurse who has subsequently brought up the child as the main carer, as in Shakespeare's *Romeo and Juliet*, where the nurse seems to be more attached to Juliet than her actual mother.

The point of the analogy is to attempt to broaden an interpretation based on close (perhaps even, blood) relationship, which might be felt not to favour anyone except parents, and to include other (non-blood) relationships, based on more specific details of the particular relationship between plaintiff and victim. Once it has

⁵*Dillon v. Legg*, 68 Cal. 2d 728 (1968). The case concerned a mother who, with her daughter, witnessed the death of her other child in a car accident and sued the driver for damages on the grounds of emotional distress. The lower court had rejected her claim because she herself was not in physical danger, but the appeal overturned that decision.

been decided that the relationship is sufficiently close, the cases become identical at the factor level: if it is not sufficiently close, the cases are similar but significantly different. Here, importantly, the analogical argument is not about similarity of the cases, but about similarity of their protagonists. Once this had been decided, factors can be assigned and the similarity or otherwise of the cases can be straightforwardly resolved.

7 ANALOGY IN A FOURTH AMENDMENT CASE

Much discussed in AI and Law are cases concerning the automobile exception to the US 4th Amendment ([44], [9], [13], [3]). The 4th Amendment prohibits searches without a warrant, subject to certain exceptions. One such exception is for automobiles (initially established in *Carroll v United States*⁶, and developed through a series of cases [13]), on the grounds that obtaining a warrant may not be possible because of the mobility of the automobile, and because automobiles have a diminished expectation of privacy, since they may be routinely stopped and examined. The case of *California v Carney*⁷ concerned a motor home⁸ which was being used to peddle marijuana while in a San Diego car park. The vehicle was searched without a warrant and marijuana was found in the rear (living) compartment of the motor home. Did the lowered expectations of privacy apply? A number of possible analogies presented themselves: it could be treated as a normal automobile since it was in a public short stay car park, and so being used as an automobile; it could be treated as a dwelling since it was designed for such use and contained a bed and a refrigerator; it could be considered akin to luggage, since it was a container of personal effects; or the rear section could be considered as akin to a glove compartment since it was an integral part of the vehicle not normally in plain sight.

In his dissent Mr Justice Stevens seemed to like the analogy with a dwelling:

“When a motor home is parked in a location that is removed from the public highway, I believe that society is prepared to recognize that the expectations of privacy within it are not unlike the expectations one has in a fixed dwelling.”

and further Mr Justice Stevens seemed to think that an analogy with luggage might also be acceptable.

“The Court in *Chadwick* specifically rejected the argument that the warrantless search was ‘reasonable’ because a footlocker has some of the mobile characteristics that support warrantless searches of automobiles. The Court recognized that ‘a person’s expectations of privacy in personal luggage are substantially greater than in an automobile,’ [433 U.S. at 433 U. S. 13], and noted that the practical problems associated with the temporary detention of a piece of luggage during the period of time necessary to obtain a warrant are significantly less than those associated with the detention of an automobile. *Id.* at 433 U. S. 13, n.7. ... It is perfectly

⁶*Carroll v. United States*, 267 U.S. 132 (1925)

⁷*California v. Carney*, 471 U.S. 386 (1985),

⁸Specifically a Dodge Mini Motor Home, one of the smaller and more basic examples of the class, resembling a van, rather than an opulent example such as a Winebago.

obvious that the citizen has a much greater expectation of privacy concerning the interior of a mobile home⁹ than of a piece of luggage such as a footlocker.”

He does, however, note the difference in that the luggage may be separately detained, which probably explains why he preferred to rely on the dwelling analogy.

The majority seemed to find the analogy with enclosed parts of an automobile more persuasive. In the majority opinion Chief Justice Burger wrote:

“But even when enclosed ‘repository’ areas have been involved, we have concluded that the lesser expectations of privacy warrant application of the exception. We have applied the exception in the context of a locked car trunk, *Cady v. Dombrowski, supra*, a sealed package in a car trunk, *Ross, supra*, a closed compartment under the dashboard, *Chambers v. Maroney, supra*, the interior of a vehicle’s upholstery, *Carroll, supra*, or sealed packages inside a covered pickup truck, *United States v. Johns*, 469 U. S. 478 (1985) ... These reduced expectations of privacy derive not from the fact that the area to be searched is in plain view, but from the pervasive regulation of vehicles capable of traveling on the public highways.”

The analogy with dwelling is rejected by Burger because

“the vehicle was so situated that an objective observer would conclude that it was being used not as a residence, but as a vehicle. ... To distinguish between respondent’s motor home and an ordinary sedan for purposes of the vehicle exception would require that we apply the exception depending upon the size of the vehicle and the quality of its appointments”

This is effectively a refusal to see the relevant aspect as a dimension running from a small car to a permanent dwelling, but as a pair of Booleans: vehicle or not vehicle and dwelling or not dwelling, with the key issue being the *actual use* of the vehicle at the time of the search. The refusal to see it as a range is justified by the treatment of luggage in previous decisions:

“we declined to distinguish between ‘worthy’ and ‘unworthy’ containers, noting that ‘the central purpose of the Fourth Amendment forecloses such a distinction.’ We decline today to distinguish between ‘worthy’ and ‘unworthy’ vehicles” (Burger).

Thus the motor home is considered by the majority to be analogous to a vehicle when parked in a downtown car park, and the rear compartment to a glove box or a boot (trunk in US). New aspects which will need to be considered in future cases are identified: the *use* to which the vehicle is being put, with *location* being an important test [1]: had the motor home been parked in a trailer park there may well have been a good case to treat it as a dwelling.

⁹Note the misleading (conscious or unconscious) use of “mobile”, rather than “motor”, home. The Wikipedia definition of the former is “a prefabricated structure, built in a factory on a permanently attached chassis before being transported to site (either by being towed or on a trailer)”, while the latter is “a type of self-propelled recreational vehicle (RV) which offers living accommodation combined with a vehicle engine.” “Motor home” was (correctly) used in the syllabus.

8 ROLES FOR ANALOGY

The use of analogy in *Dillon* is rather different from the use of analogy in *Carney*. Taken together the cases show that there are several different ways in which analogy can be used.

The new case may fall into a gap on a single aspect. *Dillon v Legg* can be seen this way, with the relevant aspect “close relation”. We may have precedents with a mother considered sufficiently close and a cousin not so considered. If a new case with a kindergarten teacher arises we may see the relationship as closer than a cousin (because love arises from caring) or more distant (because not a blood relationship). It may be held that the relationship is closer than a cousin, but not sufficiently close to a mother, because of the differences (more children, temporary relation, less intimate care). But if a further case with Juliet’s wet nurse arises, these differences do not apply and the analogy with mother may be accepted. This reasoning can be seen as determining the switching point on an aspect with magnitude.

Alternatively we may need first to choose between several possibilities for determining the aspects relevant in the case. In *Carney* we can see the relevant aspect as function (vehicle/dwelling) or container (luggage/glove box). Both the majority and the minority went for the first option, and the decision turned on the current use (as a vehicle) as against the potential use (as a dwelling). The majority held that the current use prevailed and introduced a new aspect (location) to serve as a test for current use. In *Carney* the first decision was on which aspect the analogy will be sought, and the decision was on the use of the vehicle.

Analogy can also be used in deciding on the nature of the aspect, whether Boolean or with magnitude. The majority decided, using an analogy with luggage, that this was not to be considered as a range, running from standard automobile through to standard dwelling, but as a pair of Booleans, *in-use-as-vehicle* and *in-use-as-dwelling*. Location is then used as a test to decide which applies: being on the highway or in a car park implies *in-use-as-vehicle* while being in a trailer park may imply *in-use-as-dwelling*. This clarifies that the nature of the aspects *vehicle* and *dwelling* and relates them to the third aspect, *location*.

Further, it might be that we need to consider two aspects in the analogy: in *Dillon*, Tobriner mentions the cause of the emotional distress as well as the closeness of the relationship. It may be that different causes would affect how we viewed the closeness of different relationships. This is akin to the trade offs of [16] where it is also argued that there would never be a need to consider more than two aspects simultaneously.

These points bring us back very close to the position of [51], in which Gardner distinguished between “easy cases” which could be solved by rule application and “hard cases” which required more sophisticated reasoning. Similarly [48] argues that at the top level of the argument tree rules can be used, but in order to determine whether the predicates in the antecedents of these rules apply, some kind of HYPO style case based reasoning is needed. Also in IBP [20] the top level comprises rules which identify issues, which are then resolved using CATO style case based reasoning. Both HYPO and CATO use analogical reasoning, albeit only in terms of aspects, not of elements of the cases. In neither [48] and [20], however, do the rules come from precedents. Instead they use statutes and commentaries. This

limits the extent of the rule based reasoning in those approaches, but in the approach that has emerged from this paper, rules derived from precedents can take us much further down the tree, and are sufficient to fully resolve many cases in the manner of [31]. We still, however, need analogical reasoning in certain circumstances where we have a novel situation which the current judges have to resolve. In our discussion above we found three such situations:

- when we are trying to justify a switching point (as in *Dillon*);
- when we have to choose which aspects are relevant in a case (as in *Carney*); or
- when we must determine whether an aspect is Boolean or has magnitude (also in *Carney*).

Answering these questions will, however, require something more than HYPO or CATO offer: in both those approaches the starting point for reasoning with cases is a set of aspects, but as we have seen, we may need reasoning with analogy to determine which are the applicable aspects, or even the very nature of these aspects. This need to recognise an inherent limitation in starting from a representation of cases as bundles of factors (or dimensions), which has been prevalent in AI and Law for more than 20 years, from [5] to [43], and including all the approaches outlined in section 2, is why we need to take analogy seriously. The need arises especially when the analogies are between particular real-world elements of the case and drawing attention to this possibility is one of the main motivations for the discussion in this paper.

9 PROSPECTS FOR COMPUTATIONAL MODELS OF ANALOGY

The question now arises as to whether it is feasible to attempt to provide computational support for analogical reasoning in cases presenting novel situations. Since we have moved beyond the narrow confines of law into the real world, we would need a representation of real world concepts such as vehicles and dwellings (*Carney*), mothers and teachers (*Dillon*) and baseballs and foxes (*Popov v Hayashi*, which involved a baseball, but in which analogies were drawn with *Pierson v Post*, which concerned a fox [10]). This suggests that at the heart of such support would be an ontology, which would provide the various properties of such objects which could then be mapped to produce the analogies. Such an ontology would need to be highly sophisticated: a lightweight ontology such as Wordnet [38] would not be enough, since we need to have a reasonably comprehensive list of the attributes of our objects. Moreover such an ontology would need to deal with open texture, whereas ontologies currently typically rely on definition by necessary and sufficient conditions.

In [8] Kevin Ashley discussed the requirements for an ontology which would be able to support such analogical reasoning, as well as teleological and hypothetical reasoning. Ashley's focus was on the Socratic dialogues common in US Law Schools, and what would be needed to support such discussions. His running example, explored in detail, is the wild animals cases beginning with *Pierson v Post* introduced to AI and Law in [18]. Ashley distinguishes between case-based comparisons and analogies. This is quite similar to the distinction that emerged above between analogies between cases represented as sets of factors, and analogies relating to the real-world elements of the case. Ashley claims that the former "is nearly within reach of current AI and Law technology". In the decade since [8]

was written this claim has become even more tenable. For cases seen as collections of Boolean factors (either present in or absent from a case), there is now an excellent understanding of how to compare and reason with cases, as summarised in [14]. The ontology required for this kind of reasoning can describe cases represented using factors rather than the real-world domain. Such a case related ontology can be found in [53]. To facilitate case comparison, an ontology which describes cases represented as sets of factors might well be sufficient. For example:

```
CASE:
Plaintiff: #person
Defendant: #person
PlaintiffFactors: List of #factor
DefendantFactors: List of #factor
Outcome: {plaintiff, defendant}
```

might well be enough to enable the kind of analogies drawn in [7], [5] and [52] but is entirely inadequate for the analogies discussed in sections 4 and 5 above. To support such analogies, Ashley remarked would "necessitate significant advances in the design of ontologies and the kinds of reasoning they support".

But producing such an ontology would be a colossal task, even if these advances in design had been achieved. Ashley's proposal

"is to design an ontology incrementally. The examples can be made incrementally more complex; as new features are built into the ontology, more advanced behavior can be simulated, tested, and accommodated."

This approach was explored in [1] in which a representation of the Supreme Court oral hearings is used to update an ontology for the relevant domain. In this way an ontology built on the basis of previous cases can be extended with the issues, factors and facts arising from a new case (*Carney* is used as the example in [1]).

An incremental approach can work for certain applications. The task in [1] did not involve analogical reasoning, but used oral hearing dialogues employing such reasoning to modify and extend an existing ontology to cover the case under consideration. Similarly Ashley's envisaged application is to support teaching analogy through Socratic dialogue in a law school. When preparing a course the teacher will select cases, and if she wishes to add new cases, is able to incrementally extend the ontology to cover the new cases. As the course develops over a number of years, an extensive and useful ontology may be developed, able to find and critique analogies *for that limited set of cases*. But if we wish to support the finding and critique of analogies in new cases with novel situations, this approach cannot work: such cases lead to the modification of the ontology, and so their arguments cannot be found in the existing ontology. An incrementally built ontology will inescapably be tailored to a partial set of cases in a particular domain.

In order to handle novel situations, we would need to have a comprehensive ontology already available from which to produce the analogies. In [15] Thorne McCarty gives a commentary on [8] expressing doubt that it would be worthwhile to build such an ontology: understanding what it would comprise is enough, and for practical use it would not repay the necessary investment. This is obviously true if the intended use of the ontology is to support a particular application. But what of a general purpose ontology intended to provide a comprehensive ontology and knowledge base

that would cover all the basic concepts and "rules of thumb" about how the world works? This would be intended to capture the "common sense" knowledge that people are able to bring to bear when, for example, forming and critiquing analogies in novel legal cases. The (enormous) cost of such a project might be justified by the ability to deploy such an ontology in many, many, applications (not only in law, but in every domain). It was just such an ontology that was proposed by Lenat in his CYC project which was begun in 1984 as a ten year project [37], [36]. The project has continued since 1995 under the auspices of CYCORP¹⁰. Although some successes have been reported¹¹, it has not proven to be a panacea for all knowledge base needs. It has not, to our knowledge, been tried in the context of law.

There are a number of problems with CYC which may have led to its lack of widespread use. One is the problem of completeness. No matter how large CYC becomes, it is hard to imagine that it will have everything required for every application. One of the first major uses of CYC was in the DARPA high-performance knowledge bases project [24]. Even at that time CYC was impressively large:

"The CYC knowledge base represents millions of hand-crafted axioms entered during the 13 years since CYC's inception. Through careful policing and generalizing, there are now slightly fewer than 1 million axioms in the knowledge base, interrelating roughly 50,000 atomic terms. ... Most embody general consensus information about the concepts. For example, one axiom says one cannot perform volitional actions while one sleeps, another says one cannot be in two places at once, and another says you must be at the same place as a tool to use it. The knowledge base spans human capabilities and limitations, including information on emotions, beliefs, expectations, fears, and goals; common everyday objects, processes, and situations; and the physical universe, including such phenomena as time, space, causality, and motion." [24]

None the less, CYC was not adequate to fulfill the tasks asked of it by the project, and required considerable extension to meet the specific tasks¹² on the project. In [24] we read:

"Answering the challenge problem questions takes a great deal of knowledge. For the first year's challenge problem alone, the Cycorp and Teknowledge team added some 8,000 concepts and 80,000 assertions"¹³

This might be seen as a practical problem, and that eventually CYC will be capable of use without additions. Another practical problem, however, concerns maintenance: concepts change over time. But there are some theoretical objections as well.

¹⁰<https://www.cyc.com/>

¹¹<http://www.cyc.com/enterprise-solutions/success-stories/>

¹²The project involved several "challenge problems" including *crisis management* and *battlespace integration*. For readers whose own ontology does not include *battlespace*, it is defined in the *Military Jargon Database* (<http://www.jargondatabase.com/Category/Military/Strategy-Jargon/Battle-Space>) as "The combined field of combat. For instance, an invasion of a coastal area would have a sea component (used by the Navy), a shore component (Marines), inland area (Airborne troops perhaps) and the area above it all (Air Force). The concept of the battle space is having a unified plan for all of these actors."

¹³Even the additional fragment is larger than any legal ontology yet constructed in AI and Law.

Firstly most successful ontologies have been constructed for a particular task [32]. Moreover, an ontology will often be subjective: there may be several ways of representing a group of concepts, none of which can be said to be right or wrong, even for a particular task. As [12] puts it:

"The ontology represents a conceptualisation. Different people will conceptualise a domain differently according to experience, temperament and their tasks in the domain. We should therefore not expect that there is a single ontology applicable to a domain: the point of the ontology is often to make differences in conceptualisations visible."

This last point is particularly pertinent to the ontologies in *Carney*, as discussed above. The points at dispute include:

- whether a motor home has the expectations of privacy appropriate to a vehicle or a dwelling;
- whether different vehicles have different expectations of privacy, or these expectations are equally lowered for all vehicles irrespective of their type;
- whether it is the current or potential use that determines the nature of an object (here whether the motor home is a vehicle or a dwelling).

The questions could be answered by reference to an ontology, but that would mean that we were obliged to follow the conceptualisation embodied in *that particular* ontology, whereas the case in fact turned on the difference in the conceptualisations of plaintiff and defendant. This is quite explicit in [1], in which different structures are created to represent the conceptualisations of the plaintiff, the defendant and the justices. The last of these reflects the accepted arguments, and can be seen as containing the extensions to an existing ontology required to accommodate the new case. This fits well with the idea in [8], which suggests that ontologies be constructed incrementally. The situation in cases such as *Carney* appears to be that there are competing ontologies, and the justices choose between the alternatives. In consequence, the "correct" ontology can only be determined *after* the decision. This would make any attempt to construct an ontology capable of dealing with as yet unknown novel cases futile: what such an ontology *should* contain cannot be known in advance of the case being heard, since it turns on what the majority of the judges find most plausible, when actually confronted with a particular novel case, and the accompanying arguments. The essence of open textured concepts is that decisions are made as the need arises and there is no pre-existing "truth of the matter".

9.1 Big Data and Machine Learning

These limitations of symbolic systems for predicting legal cases, together with the much increased availability of large amounts of data, have led to a growth of interest in systems using machine learning techniques to predict the outcomes of cases. Large quantities of data do exist: [33] drew on an extensive database of Supreme Court Decisions. As reported on the website of *Science* journal¹⁴ it

"used the Supreme Court Database, which contains information on cases dating back to 1791, ... They drew

¹⁴<https://www.sciencemag.org/news/2017/05/artificial-intelligence-prevails-predicting-supreme-court-decisions>

on 16 features of each vote, including the justice, the term, the issue, and the court of origin. Researchers also added other factors, such as whether oral arguments were heard ... From 1816 until 2015, the algorithm correctly predicted 70.2% of the court's 28,000 decisions and 71.9% of the justices' 240,000 votes".

Another study [22] draws on digitisations of "80,000 cases and a million judge votes from 1891 in the Circuit Courts". These projects were not chiefly concerned with predicting case outcomes, but another project which was designed for prediction and which attracted a great deal of attention was reported in [4]. It was designed to predict violations of the European Convention on Human Rights and used 584 cases. That project reported a success rate of 79%. Although these techniques undoubtedly have their uses (e.g. [22] is concerned with identifying examples of, presumably unconscious, bias in its collection of decisions), there are some doubts concerning their application to systems intended to *predict* new cases.

- *Success rate.* Typically these systems report a success rate of between 70 and 80%. While this can be seen as successful in that it does significantly improve on the baselines, it does not really represent an acceptable level of performance for use in real courts passing sentences on real people. Symbolic approaches using hand-crafted rules typically report an ability to deal with over 90% of cases (e.g. [20] and [2]). This seems reasonable, given that far fewer than 10% cases introduce the kind of challenging novel situations which we saw as giving difficulties for symbolic approaches, so it is not clear that the machine learning approaches represent an improvement in terms of prediction.
- *Transparency and Explanation.* A major strength of the symbolic approaches is that they present their reasoning as arguments which can be assessed and critiqued. In law, an answer is never enough: it is the reasons for that answer which are needed [45]. The explanation of the predictions in machine learning systems is often not satisfactory. For example [4] offered "the 20 most frequent words, listed in order of their SVM weight". These lists do not, however, look immediately promising: the list for topic 23 of article 6 of the European Convention on Human Rights predicting violation, for example, is:
 court, applicant, article, judgment, case, law, proceeding, application, government, convention, time, article convention, January, human, lodged, domestic, February, September, relevant, represented
 One doubts whether this would be the list expected by a legal expert.
- *Will the future resemble the past?* In science we generally assume that this is so. But in law we can have far less confidence. One of the strengths of the Common Law is its ability to adapt to new situations and changing social attitudes. Precedents themselves may be subject to reinterpretation in the light of subsequent decisions. Law *changes* [19], sometimes suddenly as a result of new legislation or a landmark decision, but sometimes gradually as the result of the accumulation of

numerous decisions. The latter is hard to detect, and can cast doubt on datasets of past cases¹⁵.

- *Bias.* As suggested above, some machine learning projects have identified biases in the set of past decisions. A prediction system based on such data will inevitably embody some prejudices. But enshrining such prejudices, however much it may improve prediction, is not desirable: we should be trying to eliminate these biases, not to perpetuate them.

It is tempting to look to machine learning for answers to the problems arising from the impossibility of devising the ontologies needed to support reasoning by analogy in novel cases. After all sub-symbolic techniques have succeeded in areas such as computer vision [34] when symbolic techniques failed. The above considerations, however, suggest that there are also problems - especially when trying to predict novel, "hard" cases - with machine learning approaches. We must be equally careful about the problems we address with these techniques and the claims we make for them.

10 CONCLUDING REMARKS

As noted in the review of previous approaches, precedential reasoning as rule application has become standard in AI and Law: since [40] showed how to reconstruct CATO using rules, rules have offered a way to encapsulate such reasoning formally, and in a fashion more amenable to computation. However, revisiting the topic in the light of Stevens' paper [49] we can see that a vital role remains for analogy. In Gardner's *easy cases* Stevens suggests that either rule application or analogy can be used. She does, however, argue that there is a role for analogy, especially where judicial discretion can or must be exercised, which involves interpretation of the new case in the light of precedents and possibly also reinterpretation of precedents in the light of the new case. Our discussion supports this view and we suggest that in Gardner's *hard cases*, those for which "the rules run out", it is essential. But as we have seen from the above discussion, the analogy in these cases is not between sets of factors, but between particular *elements* of the cases. In *Dillon*, the analogy is between the people who suffered trauma; in *Carney* it is between a motor home and a sedan, or the rear compartment of a motor home and a footlocker, and in *Popov v Hayashi* [10] it is between a baseball and a fox. This means that defining the relevant aspects in advance, essential for reasoning in the style of HYPO, CATO and subsequent approaches based upon them [14], is not always possible: comparing objects requires a great deal more world knowledge than comparing cases described as a set of factors (or dimensions). The analogy is not between sets of factors, but between the real world facts (which may themselves be described in a variety of ways) which determine whether or not a factor is present. Moreover, in these hard cases, the relevant facts may not have been considered in any precedent: *Carney* was the first motor home to be considered.

To satisfactorily reason with hard cases it would be necessary to model argumentation with analogy in such a way that it can handle these comparisons. It would seem that any such reasoning would need to be based on an extensive ontology, covering not only legal cases, as in [53], but also the actors and objects in these

¹⁵This criticism perhaps does not apply to [33] which was concerned with the behaviour of the justices rather than what the law should be. As such, its subject is human nature, which is less changeable than the law.

cases. But even if we ignore theoretical objections relating to task-orientation, such as subjectivity and the impossibility of choosing between alternative ontologies before a case is decided, the sheer scale of the required ontology makes it an infeasible task. While it may be possible to construct an ontology for selected past cases (perhaps supported by Natural Language Processing and Machine Learning techniques) which would be useful for reconstructing the argumentation, or for training law students in how to use and critique analogies, such an ontology would not necessarily be able to support such reasoning in new cases. It therefore appears that providing arguments for cases where the rules run out, whether to define the “switching point”, to decide whether an aspect is Boolean or has magnitude, or to determine which aspects are present in a case, must remain the province of human lawyers.

Developers of AI and Law systems have long seen their role as supporting lawyers and judges rather than replacing them. Even back in 1976, in [25], a paper which was fairly bullish on the prospects of computer judges, it was noted that Thorne McCarty, a founding father of AI and Law, had reservations about the ability of a program to “capture many of the significant [aspects] of legal reasoning”. By the mid-eighties the consensus was very much that “Computers cannot replace the lawyer, but computers can aid the legal profession in several useful ways.” [28]. Although the notion of replacing lawyers does resurface from time to time, especially in popular media¹⁶, the discussion in this paper identifies some important facets of legal reasoning which are not amenable to computer techniques, and are likely to remain so.

REFERENCES

- [1] L. Al Abdulkarim. *Representation of case law for argumentative reasoning*. PhD thesis, U of Liverpool, 2017.
- [2] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. Evaluating the use of abstract dialectical frameworks to represent case law. In *Proceedings of the 15th ICAIL*, pages 156–160. ACM, 2015.
- [3] L. Al-Abdulkarim, K. Atkinson, and T. Bench-Capon. A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
- [4] N. Aletras, D. Tsarapatsanis, D. PreoŃiu-Pietro, and V. Lampos. Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science*, 2:e93, 2016.
- [5] V. Aleven. *Teaching case-based argumentation through a model and examples*. PhD thesis, U of Pittsburgh, 1997.
- [6] L. Alexander. Constrained by precedent. *Southern Cal Law Review*, 63:1, 1989.
- [7] K. D. Ashley. *Modeling legal arguments: Reasoning with cases and hypotheticals*. MIT press, 1991.
- [8] K. D. Ashley. Ontological requirements for analogical, teleological, and hypothetical legal reasoning. In *Proceedings of the 12th ICAIL*, pages 1–10. ACM, 2009.
- [9] K. D. Ashley, C. Lynch, N. Pinkwart, and V. Aleven. A process model of legal argument with hypotheticals. In *Proceedings of JURIX 2008*, pages 1–10, 2008.
- [10] K. Atkinson. Introduction to special issue on modelling Popov v. Hayashi. *Artificial Intelligence and Law*, 20(1):1–14, 2012.
- [11] T. Bench-Capon. Arguing with cases. In *Jurix 1997*, pages 85–100, 1997.
- [12] T. Bench-Capon. Ontologies in AI and Law. *Inteligencia Artificial Aplicada al Derecho*, pages 65–85, 2005.
- [13] T. Bench-Capon. Relating values in a series of Supreme Court decisions. In *Proceedings of Jurix 2011*, pages 13–22, 2011.
- [14] T. Bench-Capon. HYPO’s legacy: introduction to the virtual special issue. *Artificial Intelligence and Law*, 25(2):205–250, 2017.
- [15] T. Bench-Capon, M. Araszkiwicz, K. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourguine, J. G. Conrad, E. Francesconi, et al. A history of AI and Law in 50 papers: 25 years of the International Conference on AI and Law. *Artificial Intelligence and Law*, 20(3):215–319, 2012.
- [16] T. Bench-Capon and K. Atkinson. Dimensions and values for legal CBR. In *Proceedings of JURIX 2017*, pages 27–32, 2017.
- [17] T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1-2):97–143, 2003.
- [18] D. H. Berman and C. D. Hafner. Representing teleological structure in case-based legal reasoning: the missing link. In *Proceedings of the 4th ICAIL*, pages 50–59. ACM, 1993.
- [19] P. Bratley, J. Frémont, E. Mackaay, and D. Poulin. Coping with change. In *Proceedings of the 3rd ICAIL*, pages 69–76. ACM, 1991.
- [20] S. Bruninghaus and K. D. Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th ICAIL*, pages 233–242. ACM, 2003.
- [21] E. Charniak and D. McDermott. *Introduction to AI*. Addison-Wesley, 1987.
- [22] D. L. Chen. Judicial analytics and the great transformation of American Law. *Artificial Intelligence and Law*, 27(1):15–42, 2019.
- [23] A. Chorley and T. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *Artificial Intelligence and Law*, 13(3-4):323–371, 2005.
- [24] P. R. Cohen, R. Schrag, E. Jones, A. Pease, A. Lin, B. Starr, D. Gunning, and M. Burke. The DARPA high-performance knowledge bases project. *AI Magazine*, 19(4):25, 1998.
- [25] A. D’Amato. Can/should computers replace judges? *Georgia Law Review*, 11:1277, 1976.
- [26] N. Duxbury. *The nature and authority of precedent*. CUP, 2008.
- [27] D. Gelbart and J. Smith. Flexicon. In *Proceedings of the 4th ICAIL*, pages 142–151. ACM, 1993.
- [28] M. M. Hazen and T. L. Hazen. Simulation of legal analysis and instruction on the computer. *Indiana Law Journal*, 59:195, 1983.
- [29] J. Horty. Reasoning with dimensions and magnitudes. In *Proceedings of the 16th ICAIL*, pages 109–118. ACM, 2017.
- [30] J. Horty. Reasoning with dimensions and magnitudes. *Artificial Intelligence and Law*, pages 1–37, Available online March 2019.
- [31] J. Horty and T. Bench-Capon. A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20(2):181–214, 2012.
- [32] D. Jones, T. Bench-Capon, and P. Visser. Methodologies for ontology development. In *XV IFIP world computer congress*, pages 62–75, 1998.
- [33] D. M. Katz, M. J. Bommarito II, and J. Blackman. A general approach for predicting the behavior of the Supreme Court of the United States. *PLoS one*, 12(4), 2017.
- [34] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [35] G. Lamond. Do precedents create rules? *Legal Theory*, 11(1):1–26, 2005.
- [36] D. B. Lenat. CYC: A large-scale investment in knowledge infrastructure. *Commun. ACM*, 38(11):33–38, 1995.
- [37] D. B. Lenat, M. Prakash, and M. Shepherd. CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI magazine*, 6(4):65, 1985.
- [38] G. A. Miller. Wordnet: a lexical database for English. *Communications of the ACM*, 38(11):39–41, 1995.
- [39] H. Prakken. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law*, 10(1-3):113–133, 2002.
- [40] H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.
- [41] E. Rich. *Artificial Intelligence*. McGraw-Hill, 1983.
- [42] A. Rigoni. An improved factor based approach to precedential constraint. *Artificial Intelligence and Law*, 23(2):133–160, 2015.
- [43] A. Rigoni. Representing dimensions within the reason model of precedent. *Artificial Intelligence and Law*, 26(1):1–22, 2018.
- [44] E. L. Rissland. Dimension-based analysis of hypotheticals from Supreme Court oral argument. In *Proceedings of the Second ICAIL*, pages 111–120, 1989.
- [45] E. L. Rissland. AI and Similarity. *IEEE Intelligent Systems*, 21(3):39–49, 2006.
- [46] E. L. Rissland and K. D. Ashley. A case-based system for Trade Secrets law. In *Proceedings of the 1st ICAIL*, pages 60–66. ACM, 1987.
- [47] A. C. Roth. *Case-based reasoning in the law: a formal theory of reasoning by case comparison*. Maastricht University, 2003.
- [48] D. B. Skalak and E. L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial intelligence and Law*, 1(1):3–44, 1992.
- [49] K. Stevens. Reasoning by precedent—between rules and analogies. *Legal Theory*, pages 1–39, 2018.
- [50] H. Turtle. Text retrieval in the legal world. *Artificial Intelligence and Law*, 3(1-2):5–54, 1995.
- [51] A. von der Lieth Gardner. *An Artificial Intelligence approach to legal reasoning*. Bradford Books, MIT Press, 1987.
- [52] A. Wyner and T. Bench-Capon. Argument schemes for legal case-based reasoning. In *Proceedings of Jurix 2007*, pages 139–149, 2007.
- [53] A. Wyner and R. Hoekstra. A legal case OWL ontology with an instantiation of Popov v. Hayashi. *Artificial Intelligence and Law*, 20(1):83–107, 2012.

¹⁶E.g. <http://www.bbc.com/future/story/20180430-are-algorithms-better-lawyers-than-humans>, <https://www.case-crunch.com/#challenge> and footnote 1.