SECTION TWO.

Towards Algorithmic Legal Reasoning and Law-Making

Computational Legal Problem Solving. What can Legal Tech Learn from AI and Law Research, and Beyond?¹

Michał Araszkiewicz

1. Introduction

LegalTech is one of the most rapidly growing branches of information technology.² It has become commonplace that advanced solutions, including natural language processing (NLP) algorithms based on machine learning (ML), may significantly increase the speed and accuracy of many juridical tasks performance. The digital transformation of international law firms and legal departments in corporations widely considers the application of tools used with problems that precede actual juridical work (paralegal tasks such as the retrieval of documents, systematization of information, and checking the formal structure of documents). It also concerns solutions that may support the assessment of similarities between legal cases, generate arguments from knowledge bases, or evaluate the relative strength of competing arguments (the tasks of lawyers). Nevertheless, computational tools to support the performance of lawyers' tasks on an effective level are difficult to develop.

There are many sources of this difficulty. Perhaps the most general observation in this connection is that many problems solved by lawyers are not well-defined. In the theory of problem solving, a problem is well defined if it has a clearly specified initial state and goal state (solution) as well as a set of operators that may be used in the transition from the initial state to the goal state.³ Among the most important specific

-
- 1 The article was financed by the National Centre for Sciences as part of research project agreement UMO-2018/29/B/HS5/01433.
 - 2 See Richard Susskind, *Tomorrow's Lawyers. An Introduction to Your Future* (Oxford University Press 2nd edn, 2017); Markus Hartung, Micha-Manuel Bues, Gernot Halbleib, *Legal Tech: How Technology Is Changing the Legal World* (C. H. Beck 2018); Jens Wagner, *Legal Tech und Legal Robots. Der Wandel im Rechtswesen durch neue Technologien und Künstliche Intelligenz* (Springer 2020).
 - 3 Kevin Dunbar, 'Problem Solving' in William Bechtel and George Graham (eds), *A Companion to Cognitive Science* (Blackwell Publishers 1999) 293–294 and a classical monograph by Allen Newell and Herbert A. Simon, *Human Problem Sol-*

issues, the following may be noted. First, it is difficult to determine the set of sources from which relevant information should be retrieved. Certain categories of sources are hardly debatable (such as statutes or, in Anglo-American legal culture, binding precedents). However, it is often unclear to what extent other sources, such as soft law or legislative materials, should be considered. Second, even if the set of relevant sources has been determined, it may be a very complex task to decide what is the structure and content of elements derived therefrom, and, in particular, how potential incompatibilities between these elements should be solved.⁴ Third, in legal reasoning, we often must decide not only based on uncertain or contradictory information, but also incomplete information. In particular, reasoning with and about evidence often involves balancing probabilities or the resolution of problems through the application of rules concerning burden of proof.⁵ As far as questions of law are concerned, the phenomenon of incompleteness is captured by the concept of legal gaps.⁶ Fourth, a major part of legal sources is expressed in natural language. Therefore, legal texts are encumbered by such well-researched phenomena as syntactic and semantic ambiguity, vagueness, context sensitivity, and open texture. Some of these phenomena are not necessarily problematic (for instance, vagueness may be effectively used in drafting legal provisions that require flexibility). In general, though, they all contribute to increased complexity.⁷ Fifth, legal reasoning often involves differences of opinion and depends on value judgments. Therefore, legal reasoning cannot be adequately reduced to a logical operation. Its adequate representation requires modeling argumentation and arguments, particularly arguments concerning values, goals, and preferences.⁸ Both the linguistic features

ving, (Prentice Hall 1972). See also Colin Lynch, Kevin D. Ashley, Niels Pinkwart, Vincent Aleven, 'Concepts, Structure and Goals: Redefining Ill-Definedness' (2009) 19 *International Journal of Artificial Intelligence in Education* 253–266.

4 cf. Carlos E. Alchourrón and Eugenio Bulygin, *Normative Systems* (Springer-Verlag 1971); Robert Alexy, *A Theory of Constitutional Rights*, transl. J. Rivers (Oxford University Press 2002); Manuel Atienza and Juan Ruiz-Manero, *A Theory of Legal Sentences* (Springer 1998).

5 See, for instance, in the context of US law: Jack Weinstein, Norman Abrams, Scott Brewer and Daniel Medwed, *Evidence* (Foundation Press 2017).

6 See, for instance Marijan Pavčnik, 'Why Discuss Gaps in the Law?' (1996) 9/1 *Ratio Juris* 72–84.

7 Timothy Endicott, *Vagueness in Law* (Oxford University Press 2000).

8 The theories of legal argumentation are presented in Eveline Feteris's *Fundamentals of Legal Argumentation. A Survey of Theories on the Justification of Judicial Decisions* (Springer 2017) and in Giorgio Bongiovanni, Gerald Postema, Antonino Rotolo,

of legal text and the general context of the legal system contribute to the complexity of legal understanding, which is why legal interpretation remains the most-investigated issue in legal theory. Sixth, humans perform legal reasoning; it obviously has (neuro)psychological grounds. The minds of lawyers are human. Therefore, they operate based on fallible heuristics and are subject to biases, decisions are made on emotional grounds and rationalized *post hoc*, etc. These and other problems are investigated under the heading of the relatively recently emerged research area law and cognitive sciences.⁹ Nonetheless, apparently, in legal practice, the only intersubjective sphere subject to evaluation is the reasoning expressed in language (in documents such as judicial opinions or lawsuits) and having a claim to rationality. It is difficult to define the rationality criteria for legal reasoning, particularly if it is our aim to develop a realistic, not idealized, model.¹⁰ Even if we assume the rationality of lawyers (as human reasoners) is limited and bounded, the relationships between the psychological (factual) and the rational in legal reasoning remain complex and not entirely clear.

Having said this, it should be noted that, irrespective of these complexities, lawyers perform their tasks on a daily basis, and often their performance is evaluated as proper or even excellent. In other words, legal experts know when a legal task is performed well, even if it is sometimes difficult to agree on the criteria of evaluation or to reveal the implicit assumptions that provide the background for the reasoning explicitly given. Decades of developing legal theory have provided imperfect, yet informative, conceptual schemes and tools that enable us to, first, analyze the phenomena contributing to the complexity of legal reasoning and, second, to reduce this complexity or help resolve problems under assumed criteria. In particular, theories of legal interpretation and legal argumentation contributed importantly to the understanding of the rational aspects of the legal decision-making process. The problems of the relationship between the intuitive and deliberative, the persuasive and the reasonable, and the

Giovanni Sartor, Chiara Valentini and Douglas Walton (eds.), *Handbook of Legal Reasoning and Argumentation* (Springer 2018).

- 9 For a recent contribution to the understanding of the relationship between law and cognitive science, see Jaap Hage, Bartosz Brożek and Nicole Vincent (eds.), *Law and Mind. A Survey of Law and the Cognitive Sciences* (Cambridge University Press 2021).
- 10 A normative model of (rational) legal reasoning is presented, for instance, in Bartosz Brożek, *Rationality and Discourse. Towards a Normative Model of Applying Law* (Wolters Kluwer 2007).

extralegal and legal features in legal decision making are the subject of vivid debates.

The question, hence, appears whether legal problems may be modeled in a manner that would justify the thesis that they are well-defined problems. A significant part of the work toward this purpose has been made in legal theory, and since the 1970s, such attempts have also been made in the broad research area called artificial intelligence and law (AI and law).¹¹ Basically, research in AI and law consists of the use of artificial intelligence (AI) tools to build models of legal reasoning and other systems that may support the performance of juridical tasks. To characterize this area of research more accurately, we must address the scope and nature of general AI research. For obvious reasons, we cannot enter into the complex philosophical debate concerning the notion of AI. For the purposes of this paper, it is sufficient to recall the recent definition provided by one of the most prominent researchers in the field, according to whom AI pursues the goal of creating intelligent machines. A machine may be considered intelligent if it “chooses the actions that are expected to achieve its objectives, given what is perceived.”¹² The objectives are, of course, not a machine’s own, but objectives specified by the developer or reconstructed by the machine based on initial specifications. In other words, an intelligent machine should act under the principles of instrumental rationality regarding a certain set of objectives. A machine may be developed to simulate a human’s thinking or behavior or to surpass human capabilities and performance to realize said objectives.¹³

There are two broad and internally differentiated streams of research in AI: symbolic AI and computational intelligence. Although the division between these streams is not rigid, and the classification of some system types into the categories is subject to debate, they may be juxtaposed, in a model account, in the following manner:¹⁴

In symbolic AI, the model of an intelligent system is typically represented explicitly. The data used by the model have a symbolic character;

11 For a history of this scientific movement, see Trevor Bench-Capon and others ‘A History of AI and Law in 50 Papers: 25 Years of the International Conference on AI and Law’ (2012) 20 *Artificial Intelligence and Law* 215–319.

12 Stuart J. Russell, “Artificial Intelligence. A Binary Approach,” in S. Matthew Liao (ed), *Ethics of Artificial Intelligence* (Oxford University Press 2020) 327.

13 Stuart Russell and Peter Norvig, *Artificial Intelligence. A Modern Approach*, (3rd ed. Pearson 2016) 2.

14 Cf. M. Flasiński, *Introduction to Artificial Intelligence* (Springer International Publishing 2016) 15 and 23.

they are expressed in each formal language, such as the language of logic, graphs, or set theory. The operations performed by the system consist of formal operations on sets of symbols; for instance, they may have the character of deductive operations. Exemplary approaches represented in this stream are logic-based systems, rule-based systems, case-based reasoning systems, argumentation systems, and systems based on the Semantic Web architecture.

In computational intelligence, the model of an intelligent system (and the knowledge used therein) typically has implicit character. The basic type of data used by the system is numeric data. The operations performed by the system have a primarily mathematical character. The types of systems developed in this field are very diverse. For instance, we may mention support vector machines, neural networks, and evolutionary algorithms. However, the model characteristics of these models are generally applicable to all of them. In some cases, computational intelligence systems have a distributed character in that meaning cannot be ascribed to the elements of the system, but is, rather, inferred from the operations of the total system or a reasonable part thereof.

A common feature of computational intelligence mechanisms is that they can learn. The field of computational intelligence should not be identified with machine learning (ML), because there are also symbolic learning models, and not all computational intelligence models have this feature. However, most successfully applied ML systems are based on computational intelligence.¹⁵ ML is one of the most rapidly developing fields of computer science. We distinguish three main categories of ML: supervised, unsupervised, and reinforcement learning (There are also intermediary categories and finer-grained distinctions.) Generally speaking, the idea of ML is that an algorithm gradually corrects its performance to achieve the expected or desired result.

During the first decades of research on AI (1950s–1980s), the symbolic approach was dominant. This led to the development of tools and approaches applicable to many areas. In the 1970s, the concept of expert systems, that is, the systems that simulate the reasoning of domain experts, rose to prominence. Frequently, expert systems are based on rules (understood here as conditional clauses of the form “IF condition THEN

15 See Ethem Alpaydin, *Machine Learning, The New AI* (The MIT Press 2016); Miroslav Kubat, *An Introduction to Machine Learning* (Springer International Publishing 2017).

action”).¹⁶ Such systems may perform both forward-chaining reasoning (inferring conclusions from a given set of facts that may make the rules fire) and backward-chaining reasoning (verification of hypotheses posed by the user regarding the system’s knowledge). The development of case-based reasoning systems¹⁷ and defeasible logics was initiated in the 1980s.¹⁸ Subsequently, the development of argumentation systems from the 1990s on¹⁹ led to the further progress of symbolic AI. Yet another direction of development was the models of structured knowledge: In this approach, engineers intended to represent knowledge not only on the level of logical formulae, but also to consider the internal structure of concepts and connections between them.²⁰ As is known, this approach was extended and elaborated in the Semantic Web—a collection of standards enabling the presentation of semantic knowledge in machine-readable form.²¹

Symbolic AI systems, despite the many differences between them, share certain strong and weak features. Their important advantage is their high degree of understandability for the user. In principle, symbolic AI systems may present an *explanation* of the reasoning process, both regarding the general model and particular reasoning. Moreover, the reasoning they perform, based on deductive or quasi-deductive reasoning patterns, is very reliable, although, of course, the generated conclusions concern the adopted set of premises, which may be erroneous or doubtful. Conversely, symbolic AI systems may only perform based on the knowledge stored in their knowledge base or explicitly provided by the user. The knowledge

-
- 16 A very influential position on this subject is that of Bruce Buchanan, Edward H. Shortliffe, *Rule-based Expert Systems. The MYCIN Experiments of the Stanford Heuristic Programming Project* (Reading 1984).
 - 17 For more recent elaborations of this topic, see David B. Leake, ‘Case-Based Reasoning,’ in William Bechtel and George Graham (ed), *A Companion to Cognitive Science* (Blackwell Publishers 1999) 465–476; Michael M. Richter, Rosina O. Weber, *Case-Based Reasoning. A Textbook* (Springer-Verlag 2013).
 - 18 Ray Reiter, ‘A Logic for Default Reasoning,’ (1980) 13 *Artificial Intelligence* 81–132; John L. Pollock, ‘Defeasible Reasoning’ (1987), 11 *Cognitive Science* 481–518.
 - 19 Phan Minh Dung, ‘On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-person Games’ (1995) 77(2) *Artificial intelligence* 321–357.
 - 20 Allan M. Collins and Ross M. Quillian, ‘Retrieval Time from Semantic Memory’ (1969) 8(2) *Journal of Verbal Learning & Verbal Behavior*, 240–247; Marvin Minsky, ‘A Framework for Representing Knowledge’ in Patrick H. Winston (ed) *Psychology of Computer Vision* (MIT Press 1975).
 - 21 Liyang Liu, *Introduction to the Semantic Web and Semantic Web Services* (Chapman and Hall/CRC 2019).

base itself must be formalized, validated, and maintained, which are time-consuming and costly processes. Symbolic systems' capacity to learn and adjust their behavior to new situations is very limited. Moreover, it became apparent decades ago that, to perform in a satisfactory manner, many types of symbolic AI systems must be equipped with and process common-sense knowledge. However, the amount of common-sense knowledge required to attain the desired performance results of the systems is immensely great.²² Problems regarding the preparation of relevant knowledge bases for symbolic AI systems are discussed under the heading of the *knowledge acquisition bottleneck*. These disadvantages have caused the limited applicability of symbolic AI systems, which typically solve problems in narrow, well-defined domains.

The main advantages of computational intelligence systems come from their ability to learn based on (numeric) data. The emergence of Internet technologies in the 1990s led to the creation of big data sets, which enabled the spectacular development of ML techniques and tools. Nowadays, applications based on computational intelligence are naturally called *AI* because of their capacity to adapt to the context and changing circumstances, their relative autonomy and the degree of unpredictability following from it, and their high-level performance of tasks in numerous areas including scientific discovery, finances, insurance, transportation, military applications, medicine, automated translation, or the ability to conduct a natural conversation, where the latter areas exemplify the successful application of natural language processing (NLP) technologies.²³ The success of these systems follows from many factors, including their overall performance level (high speed and accuracy) and the possibility of the ongoing development of the systems throughout the learning process. Moreover, appropriately trained ML systems generally do not require the preparation of the data they operate on; some may perform well based on raw data. The ubiquitous character of solutions based on ML models has led to questions concerning the ethical and legal consequences of

-
- 22 The most famous project developing the fullest possible base of commonsense knowledge is CYC, initiated in 1984. At first, it was assumed the task would be completed by the 1990s. As of now, the CYC database contains 25 million assertions and is still growing. See <<https://cyc.com/>> access 8 May 2021.
- 23 On the topic of NLP, see Nitin Indurkha, Fred J. Damerau, *Handbook of Natural Language Processing* (Chapman & Hall/CRC 2010). The applications and models of ML systems are discussed in a non-technical manner by Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (Brilliance Audio 2017).

their operation.²⁴ Accordingly, one of the most debated issues here is the relatively low level of these systems' explainability.²⁵ Although engineers understand principles concerning the operations of the models, in many cases, it is practically impossible to state why, in a concrete situation, the algorithm generated a given result. Moreover, if the system is used in the context of decision-making support, even obtaining a detailed explanation of the algorithm's operation cannot count as adequate justification for the decision because of its distinct nature—computational intelligence performs tasks through arithmetic operations, not symbolic reasoning. The use of ML systems also generates the risk of algorithmic bias affecting the results.²⁶

Having briefly characterized the main approaches in the field of AI research, let us discuss the most important achievements of AI and law research, beginning with symbolic models of legal reasoning and continuing with a comment on ML applications and concepts combining the two approaches. In the final part of the chapter, we will discuss the basic conclusions that follow from these analyses for the LegalTech industry.

2. Modeling Legal Reasoning and Argumentation

How are legal conclusions generated and justified based on available knowledge in a computational system? As discussed above, an answer to this question requires a more precise definition of the problem being solved. The history of AI and law represents the evolution of perspectives concerning the nature of the problem solved regarding the generation and justification of legal conclusions. In many important respects, this evolution mirrors the development of the theories of legal reasoning elaborated

-
- 24 These issues have recently been discussed, *inter alia*, in a contributed volume: Maria-Jesús González-Espejo and Juan Pavón (eds.), *An Introductory Guide to Artificial Intelligence for Legal Professionals*, (Kluwer Law International 2020).
- 25 The notion of explainability and related concepts are discussed thoroughly in Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bénézet, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-Lopez, Daniel Molina, Richard Benjamins, Raja Chatila and Francisco Herrera, 'Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI' (2020) 58 *Information Fusion* 82-115.
- 26 Philipp Hacker, 'Towards a Flexible Framework for Algorithmic Fairness', in: Ralf H. Reussner, Anne Koziol, Robert Heinrich (eds.) *50. Jahrestagung der Gesellschaft für Informatik, INFORMATIK 2020 – Back to the Future* (Karlsruhe 2020) 99-108.

in legal theory. Nevertheless, it should be emphasized that the influence of legal-theoretical work on AI and law has been rather limited.²⁷

Considering the ordering based on historical precedence and the increasing complexity of the proposed models, the initial approach taken in AI and law concerned the representation of legal reasoning in rule-based systems. This view is generally based on the syllogistic model of the application of law and has been present in various legal expert systems and in models of statutory legal reasoning. The structure of legal knowledge is represented as a set of rules understood as conditional expressions of the form “IF ... THEN ...” or similar. The model of reasoning has often been implemented as a logic program. Therefore, the solution to a legal problem is defined as creating logical proof from a set of premises to a conclusion or verifying whether a conclusion is provable based on the premises.

The interface of a rule-based legal expert system typically enables the user to enter information as answers to the questions asked by the program: yes/no questions or questions about numeric information, such as the date of a certain event or some person’s age. Based on information provided (called *facts*), the program inferred conclusions. If a given conclusion was not provable based on available information, the program could infer a negative answer if it used the “negation as failure” solution. Importantly,

27 Sometimes even more far-reaching claims are made, for instance Thomas F. Gordon in the presentation made at the ICAIL 2007 conference remarks that “legal philosophy failed to provide the necessary theoretical foundation for our field” (referring to the AI and Law research), Thomas F. Gordon, ‘20 Years of ICAIL – Reflections on the field of AI and Law’, 2007, <http://www.tfgordon.de/publications/> (access 10 May 2021). The causes of this limited flow of information between legal theory and AI and Law research require thorough, systematic investigation. However, there are also examples of fruitful adoption of legal-theoretical frameworks in formal and computational models, as in Jaap Hage, ‘Formalizing legal coherence’ in Ronald Prescott Loui (ed) *Proceedings of the Eighth International Conference on Artificial Intelligence and Law*, ICAIL 2001 (ACM 2001) 22-31; Kevin D. Ashley, ‘An AI model of case-based legal argument from a jurisprudential viewpoint’ (2002) 10 *Artificial Intelligence and Law* 163-218; Giovanni Sartor, ‘Doing justice to rights and values: teleological reasoning and proportionality’ (2010) 18(2) *Artificial Intelligence and Law* 175-215 or John Henderson and Trevor Bench-Capon, ‘Describing the Development of Case Law’ in Floris Bex (ed) *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, ICAIL 2019 (ACM 2019) 32-41.

systems of this type were often accompanied by an explanatory module that presented the reasoning for the program step by step.²⁸

The main disadvantages of classical rule-based models of legal reasoning are as follows. They do not capture the dialectic features of legal reasoning, which often involves comparing arguments and balancing interests. Their linear account of reasoning cannot represent these aspects. Moreover, they require that the facts of cases introduced by the user be expressed in terms already used in the rules base stored by the system. This is an unrealistic feature of these models, because in real-life situations, legal cases typically are not directly describable in the highly general language of legal rules. Therefore, rule-based legal expert systems require the user to decide whether a particular legal category applies to a given case—where not only an unqualified user but also a professional lawyer may have doubts. This is particularly visible concerning the applicability of vague or open-textured legal concepts. The meaning of such concepts is typically subject to evolution in case law.

The rule-based approach to the modeling of legal reasoning was contested in the 1980s by the proponents of another approach: case-based reasoning models. In Anglo-American legal culture, the essence of legal reasoning seems to be captured in reasoning about the applicability of precedent cases to current factual situations and in arguing about the similarities and dissimilarities of cases. Nowadays, case-based reasoning is one of the most important areas in domains of AI research.²⁹ The paradigm of case-based reasoning modeling in AI and law was created in connection with the development of the HYPO system by Kevin D. Ashley and Edwina Rissland.³⁰ The program uses the knowledge base of cases

28 The classical contributions representing this approach are Donald A. Waterman and Mark A. Peterson, 'Models of Legal Decision Making: Research Design and Methods', (Rand Corporation, The Institute for Civil Justice 1981) and Sergot (n 208) 370–386.

29 Richter, Weber, (n 17).

30 The most complete presentation of HYPO is to be found in the monograph Kevin D. Ashley, *Modeling Legal Argument. Reasoning with Cases and Hypotheticals* (MIT Press 1990). Other accounts of case-based reasoning are also present in the literature, as in the model based on the notion of prototype and its deformations, see L. Thorne McCarty, 'An Implementation of Eisner v. Macomber' in L. Thorne McCarty (ed) *Proceedings of the Fifth International Conference on Artificial Intelligence and Law, ICAIL '95* (ACM 1995) 276–286, or in the model based on the so-called exemplar-based explanations, see L. Karl Branting, 'Building explanations from rules and structured cases' (1991) 34(6) *International Journal of Man-Machine Studies* 797–837.

(in the domain of trade secret law) indexed by dimensions—knowledge representation tools representing a scale from the most pro-plaintiff to the most pro-defendant point. Dimensions represent ordered sets of general aspects of the case, and they form the foundation of the construction of arguments based on similarities and dissimilarities between the case at bar and the quoted cases. Notably, instead of suggesting one possible answer, HYPO generated a three-ply argument naturally representing the exchange of positions in the litigation process: the first ply by the plaintiff, a reply by the defendant, and a rebuttal by the plaintiff. HYPO generated arguments based on similar cases and distinguishing arguments, as well as arguments based on counterexamples. It also pointed out the hypothetical variations of the analyzed cases to show how the argumentation for each side could be strengthened.

Numerous computational models of legal reasoning were based on the basic ideas expressed in the HYPO model or developed in directions.³¹ A particularly influential approach was proposed in CATO—a program developed to support legal education.³² In CATO, the cases were characterized by binary factors as opposed to scalable dimensions. A factor may be either present or absent in the description of the case, and if it is present, it always favors a decision for the same side (defendant or plaintiff). CATO ordered factors into a hierarchy, going from base-level factors to abstract factors, connected by positive or negative links of strength. It generated argument structures like HYPO. However, it was based on a more extensive case set, yet it still comprised the same domain of law (trade secrets law).

Other developmental directions of the computational models of legal reasoning were as follows. Rule-based and case-based reasoning were combined in hybrid systems, where case law served as the basis for establishing semantics of rules' conditions.³³ The factor-based approach was soon supplemented by teleological considerations and led to the development of the systems, which represent not only arguments about similarities or

31 Recently, the evolution of this paradigm was summarized in Trevor Bench-Capon, 'HYPO'S legacy: introduction to the virtual special issue' (2017), 25(2) *Artificial Intelligence and Law* 205-250.

32 Vincent Alevan, *Teaching Case-Based Argumentation Through A Model and Examples*, (University of Pittsburgh 1997) <<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.47.3347&rep=rep1&type=pdf>> access 10 May 2021.

33 Edwina L. Rissland, David B. Skalak, 'CABARET: Rule Interpretation in a Hybrid Architecture' (1991) 34(6) *International Journal of Man-Machine Studies* 839-887.

dissimilarities between cases, but also legal values and goals.³⁴ The factor-based approach was combined with the research on nonmonotonic logic, which led to the concept of representing legal cases as rules that connect a collection of factors (representing legally relevant information about the case's circumstances) and the case's outcome.³⁵ This latter approach has become particularly influential and led to the formalization of the models of case-based reasoning.³⁶ However, it is debated whether it represents the specific features of analogical reasoning that should be distinguished from typical rule-based reasoning.³⁷

In summing up the above considerations, it should be observed that in the dimensions- and factor-based systems of legal reasoning the task of finding and justifying a solution to a case is defined as selecting the outcome which has the strongest support with regard to the existing case base. The case-based reasoning systems does not have to yield an unequivocal answer – some of them provide a set of arguments pro and contra without determining the final solution.

The direction of research that considered types of legal arguments led to the generalized view consisting of the representation of legal reasoning as argumentation. Even though theories of legal argumentation have been discussed in the legal literature since the 1950s,³⁸ this approach to the computational modeling of legal reasoning has become dominant in the 1990s in connection with the emergence of a new paradigm for representing ar-

34 This discussion was initiated in AI and Law by Donald H. Berman and Carole D. Hafner in the paper titled 'Representing Teleological Structure in Case-based Legal Reasoning: The Missing Link' in Anja Oskamp and Kevin Ashley (eds.), *Proceedings of the Fourth International Conference on Artificial Intelligence and Law, ICAIL '93*, (ACM 1993) 50-59.

35 Henry Prakken, Giovanni Sartor, 'Modelling Reasoning with Precedents in a Formal Dialogue Game' (1998) 6(2-4) *Artificial Intelligence and Law* 231-287.

36 For instance, John F. Horty, 'Reasoning with dimension and magnitudes' (2019) *Artificial Intelligence and Law* 27(3), 309-345 and Henry Prakken, 'Comparing Alternative Factor- and Precedent-Based Accounts of Precedential Constraint', in: Michał Araszkiewicz, Víctor Rodríguez-Doncel (eds.): *Legal Knowledge and Information Systems - JURIX 2019: The Thirty-second Annual Conference*, *Frontiers in Artificial Intelligence and Applications* 322, (IOS Press 2019), 73-82.

37 Katie Atkinson and Trevor Bench-Capon, 'Reasoning with Legal Cases: Analogy or Rule Application?', in: Floris Bex (ed) *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, ICAIL 2019* (ACM 2019), 12-21.

38 Stephen Toulmin, *The Uses of Argument*, Cambridge University Press 2003 (1st ed. 1958); Chaim Perelman, Lucie Olbrechts-Tyteca, *The New Rhetoric. A Treatise on Argumentation*, (University of Notre Dame Press 1971) (originally published in French in 1958).

gumentation in formal systems, namely, argumentation frameworks.³⁹ We distinguish between abstract and structured argumentation frameworks. In abstract argumentation frameworks, the notions of argument and attack between arguments remain undefined, but they are sufficient to define the criteria for argument acceptability (so-called semantics), which generate extensions, intuitively, sets of arguments that can be rationally accepted together. Structured argumentation frameworks enable the presentation of relationships between premises and conclusions of arguments, as well as types of attacks on arguments: undermining attacks directed against arguments' premises, rebuttal attacks that question conclusions, and undercutting attacks that aim at weakening the relationship between the premises and the conclusion.⁴⁰ The computational models of argumentation as abstract structures enabling the representation of any type of argument have become extremely influential in AI and law.⁴¹ Certain aspects of legal argumentation have received their computational representation (not necessarily based on the concept of argumentation frameworks), such as reasoning with standards of proof in the Carneades system,⁴² balancing reasons in reason-based logic,⁴³ or, more recently, reasoning with burden of persuasion in a model based on ASPIC+.⁴⁴

To some extent, another approach was developed. As noted, legal reasoning may be represented not as a "battle of arguments" but rather as

-
- 39 Dung (n 244). See the general elaboration of the topic of formal and computational argumentation in Iyad Rawhan, Guillermo R. Simari (eds.), *Argumentation in Artificial Intelligence*, (Springer 2009) and in Pietro Baroni, Dov Gabbay, Xavier Parent, Leon van der Torre (eds.) *Handbook of Formal Argumentation*, (College Publications 2018).
 - 40 See Henry Prakken, 'An abstract framework for argumentation with structured arguments' (2010) *Argument and Computation* 1(2), 93-124.
 - 41 The influence of Dungian argumentation frameworks on the AI and Law research is discussed by Trevor Bench-Capon, 'Before and after Dung: Argumentation in AI and Law', *Argument and Computation* 11(1-2), 221-238.
 - 42 Thomas F. Gordon, Henry Prakken, Douglas Walton, 'The Carneades model of argument and burden of proof' (2007) *Artificial Intelligence* 171(10-15), 875-896.
 - 43 Jaap C. Hage, *Reasoning with Rules. An Essay in Legal Reasoning and its Underlying Logic*, (Springer 1997).
 - 44 Roberta Calegari and Giovanni Sartor, 'A Model for the Burden of Persuasion in Argumentation', in: Serena Villata, Jakub Harašta and Petr Kremen (eds) *Legal Knowledge and Information Systems - JURIX 2020: The Thirty-third Annual Conference*, *Frontiers in Artificial Intelligence and Applications* 334, (IOS Press 2020), 13-22; an introduction to the ASPIC+ system: Sanjay Modgil, Henry Prakken, The ASPIC+ framework for structured argumentation: a tutorial, (2014) *Argument and Computation* 5(1), 31-62.

an endeavor to construct the most coherent set of elements (theory) that explains the solution of the case.⁴⁵ Models of legal reasoning based on the notion of coherence were developed earlier in legal theory,⁴⁶ and they influenced the coherentist approach in AI and law to a limited extent. In the computational modeling of legal reasoning, coherence-based models have been most intensively investigated in the context of case-based reasoning systems, combining reasoning based on rules, factors, and values, and introducing the external criteria enabling the comparison of theories.⁴⁷ The coherence-based approach has also been successfully combined with the argument-based approach in a hybrid theory of reasoning with evidence.⁴⁸

Much attention has also been devoted to the problems of legal knowledge representation. As is known, logical formalisms, such as first-order logic or deontic logic, have limited expressive power, and they cannot account for the complex structure of legal concepts. In this connection, the ideas elaborated in the general research on AI in structured knowledge modeling have been applied widely in the fields of AI and law.⁴⁹ For instance, a frame-based approach to knowledge, representing entities as sets of slots that give information on the values of the parameters of this object, has been applied to the representation of legislation.⁵⁰ The development of Semantic Web technology has had a definitive impact on the representation of knowledge in the fields of AI and law.⁵¹ For the sake of recall, the

45 See L. Thorne McCarty, 'Some Arguments About Legal Arguments', in John Zeleznikow, Daniel Hunter, L. Karl Branting (eds.): *Proceedings of the Sixth International Conference on Artificial Intelligence and Law, ICAIL '97*, (ACM 1997), 215-224.

46 For instance Aleksander Peczenik, *On Law and Reason*, (Springer 2008).

47 See Trevor Bench-Capon, Giovanni Sartor, 'A model of legal reasoning with cases incorporating theories and values', *Artificial Intelligence* 150(102), 97-143. An alternative approach based on constraint satisfaction conception of coherence as outlined by Paul Thagard, *Coherence in Thought and Action*, (The MIT Press 2000), was applied to the field of legal reasoning in Michał Araszkiewicz, 'Limits of Constraint Satisfaction Theory of Coherence as a Theory of (Legal) Reasoning' in Michał Araszkiewicz and Jaromír Šavelka (eds) *Coherence. Insights from Philosophy, Jurisprudence and Artificial Intelligence* (Springer 2013), 217-241.

48 Floris Bex, *Arguments, Stories and Criminal Evidence. A Formal Hybrid Theory*, (Springer 2011).

49 See Erich Schweighofer, *Legal Knowledge Representation*, (Kluwer Law International 1999).

50 See Robert van Kralingen, *Frame-based Conceptual Models of Statute Law*, (Kluwer Law International 1995).

51 See for instance Pompeu Casanovas, Monica Palmirani, Silvio Peroni, Tom van Engers and Fabio Vitali 'Special Issue on the Semantic Web for the Legal Domain

Semantic Web is a multi-layered system of information that aims at facilitating the processing of information by machines. An important part of this framework is provided by ontologies—formal and computational representations of the relationships between concepts, and reasoners—computer programs that perform inferences based on information stored in an ontology.⁵² Numerous legal ontologies have been developed since the 1990s, including sophisticated ontologies of causal links⁵³ or systems representing types of legal provisions and deontic modalities.⁵⁴ The research on legal ontologies has important connections with legal-theoretical research on legal concepts.⁵⁵

In recent years, much attention has been devoted to a more natural representation of legal arguments in computational systems. Doug Walton's philosophical conception of argumentation schemes has been applied in numerous domains of AI and law research, most recently in connection with interpretive argumentation.⁵⁶ The topic of legal interpretation has

Guest Editors' Editorial: The Next Step' (2016) *Semantic Web Journal* <<http://www.semantic-web-journal.net/content/special-issue-semantic-web-legal-domain-guest-editors%E2%80%99-editorial-next-step>> access: 16 August 2021.

- 52 Nuria Casellas, *Legal Ontology Engineering. Methodologies, Modelling Trends, and the Ontology of Professional Judicial Knowledge*, (Springer 2011); Giovanni Sartor, Pompeu Casanovas, Maria Angela Biasiotti, Meritzell Fernández-Barrera (eds.), *Approaches to Legal Ontologies. Theories, Domains, Methodologies*, (Springer 2011); Johannes Scharf, *Künstliche Intelligenz und Recht. Von den Wissensrepräsentation zur automatisierten Entscheidungsfindung*, (Weblaw 2015).
- 53 Jos Lehmann and Aldo Gangemi, 'An ontology of physical causation as a basis for assessing causation in fact and attributing legal responsibility' (2007) 15(3) *Artificial Intelligence and Law*, 301-321.
- 54 For instance Enrico Francesconi, 'A description logic framework for advanced accessing and reasoning over normative provision' (2014) *Artificial Intelligence and Law* 22(3), 291-311.
- 55 Giovanni Sartor, 'Legal concepts as inferential nodes and ontological categories' (2009) *Artificial Intelligence and Law* 17(3), 217-251.
- 56 The most comprehensive presentation of the argumentation schemes theory is Douglas Walton, Chris Reed and Fabrizio Macagno, *Argumentation Schemes*, (Cambridge University Press 2008). The application of this theory to interpretive argumentation in law may be found in the recent monograph by Douglas Walton, Giovanni Sartor and Fabrizio Macagno, *Statutory Interpretation: Pragmatics and argumentation*, (Cambridge University Press 2020). The influence of Douglas Walton's theories on AI and Law has recently been discussed in Katie Atkinson, Trevor Bench-Capon, Floris Bex, Thomas F. Gordon, Henry Prakken, Giovanni Sartor, Bart Verheij, 'In memoriam Douglas N. Walton: the influence of Doug Walton on AI and law' (2020) *Artificial Intelligence and Law* 28(3), 281-326 and in Katie Atkinson and Trevor Bench-Capon, *Argumentation Schemes in AI and Law* (in press 2021).

become one of the most intensively debated issues in AI and law, including the strategies of agents performing the interpretation, the types of conflicts between interpretive statements, and the formal representation of interpretive disagreement in argumentation frameworks.⁵⁷ Generally speaking, the representation of legal reasoning in argumentation systems assumes that the correct solution is the one determined by the adopted argument acceptance criteria.

The outline of approaches present in computational models of legal reasoning (including the models of argumentation) indicates the increasing complexity of the developed approaches, as well as increasing awareness of the complexities of legal reasoning in AI and law research, even though it must be stressed again that the flow of information between this field of research and legal theory remains rather limited. There is also a visible tension between the direction focused on more informal, natural, descriptively adequate modeling (for instance, based on argumentation schemes) and strictly formal, computationally oriented modeling (as in abstract argumentation frameworks). Moreover, there is an apparent tendency to develop general, formal models as opposed to domain-dependent models that rely primarily on juridical knowledge. The fundamental question, then, emerges: Should computational models of legal reasoning simulate the bounded rationality of legal decision makers and arguers? Alternatively, should it represent legal reasoning as it would be performed by an idealized, rational entity? This problem also has a bearing on the modeling of legal prediction tasks, as discussed in what follows.

57 See for instance Michał Araszkiewicz, 'Towards Systematic Research on Statutory Interpretation in AI and Law', in: Kevin D. Ashley (ed.) *Legal Knowledge and Information Systems - JURIX 2013: The Twenty-Sixth Annual Conference*. Frontiers in Artificial Intelligence and Applications 259, (IOS Press 2013), 15-24; Tomasz Żurek and Michał Araszkiewicz, 'Modeling teleological interpretation' in Enrico Francesconi and Bart Verheij (eds), *International Conference on Artificial Intelligence and Law, ICAIL '13*, (ACM 2013), 160-168; Michał Araszkiewicz and Tomasz Żurek, 'Interpreting Agents' in Floris Bex and Serena Villata (eds), *Legal Knowledge and Information Systems - JURIX 2016: The Twenty-Ninth Annual Conference*. Frontiers in Artificial Intelligence and Applications 294 (IOS Press 2016) 13-22; Martín O. Moguillansky, Antonino Rotolo, Guillermo Ricardo Simari, 'Hypotheses and their dynamics in legal argumentation' (2019) Expert Systems and Applications 129, 37-55. General models of formal argumentation which share basic ideas of the argumentation schemes theory may be found in Bart Verheij, 'DefLog: on the Logical Interpretation of Prima Facie Justified Assumptions', (2003) 13 (3) Journal of Logic and Computation 319-34 and in Bart Verheij 'Artificial Argument Assistants for Defeasible Argumentation', (2003) 150 (1-2) Artificial Intelligence 291-324.

Notwithstanding all the important differences between the presented approaches, they all share a feature: to operate, they need a formalized knowledge base to be prepared, validated, and maintained. These processes are costly and time consuming, and they require a degree of debatable, sometimes arbitrary, decisions concerning the formalization of knowledge elements. Important choices have also been made respecting the selection of inference mechanisms performed by a system. They cannot generalize the available knowledge or analyze data that is not represented in each formal language. Therefore, the scope of their application is severely limited. What is more, as with any symbolic AI system, they may require the use of commonsense knowledge, which is the problem discussed above in the context of general AI. The advantage of systems of this type is that they can provide the reasons for the generation of a conclusion in a manner that is, in principle, understandable for a user. These reasons may have different structures, considering the diverse architectures of the systems. For instance, in classical rule-based systems, the reasons will be presented as premises of logical inference; in case-based reasoning systems, as dimension- or factor-based similarities or dissimilarities providing a basis for arguments; and in coherence-based systems, as the degrees of coherence of theories supporting given conclusions. In recent literature, it has been claimed that these systems basically generate explanations of legal decisions.⁵⁸ More strictly speaking, they generate justificatory reasoning (argumentation) in the first place, although many also can explain why and how such and such justificatory reasoning was generated. Despite these advantages, symbolic AI systems are not widespread in practice due to their low scalability, lack of possibility of analyzing source documents, and very limited adaptive capability. It should also be emphasized that, from the perspective of legal theory and doctrine, the computational models of legal reasoning may be assessed as too simplified on the conceptual level.

3. Computational Intelligence for Legal Tasks: How to Combine it with Symbolic Legal Reasoning Models

Although computational intelligence models, including neural networks, have been investigated in connection with solving legal problems since

58 See Katie Atkinson, Trevor Bench-Capon and Danushka Bollegala, 'Explanation in AI and law: Past, present and future' (2020) *Artificial Intelligence* 289: 103387.

the 1980s,⁵⁹ it was the unprecedented development of ML technology in the 21st century that established this approach as a dominant stream in AI and law. Nowadays, a significant part of research in AI and law focuses on developing systems aiming at the prediction of judicial decisions, classification of elements of legal texts, extraction of information from datasets, e-discovery, or enhancing the performance of retrieval systems. A substantial part of this research is based on NLP technology. The advances of this approach have been enabled by large datasets of legal documents available online.

The function of ML models is to identify a pattern in the dataset, considering the patterns of data already identified.⁶⁰ The general methodology for developing ML models in the field of law may be characterized as follows, taking the supervised learning approach as an example. The first step consists of the identification and gathering of a set of raw data (for instance, textual documents, such as judicial decisions). The next steps concern the preparation of the dataset, which encompasses normalization, tokenization, and annotation.⁶¹ *Normalization* consists of converting all words to lower case and eliminating variations, such as conjugation. *Tokenization* involves the elimination of punctuation or hyphens and results in the treatment of certain words or sets of neighboring words as tokens (n-grams). *Annotation* consists of adding information to the source by labeling the parts thereof. These labels may indicate the grammatical function of expressions, disambiguate them, or indicate the nature of semantic or other information carried by them (for instance, if the aim of the model is to extract the argument elements from a judicial opinion, the annotation categories may be the premises of arguments, their conclusions, and the names, or types, of arguments employed by judges). Annotation may be applied to levels of granularity; it may relate to the whole document, to parts thereof, or to phrases or words. Both the elaboration of the annotation scheme and the very process of annotation require the adoption of certain principles and the resolution of differences of opinion. In many cases, annotators eventually make decisions; the scope of convergence between them is measured under the heading of interannotator agreement.

59 See Richard K. Belew, 'A connectionist approach to conceptual information retrieval' (Proceedings of the First International Conference on Artificial Intelligence and Law, ICAIL '87, Boston, 27-29 May 1987) 116-126.

60 See Kevin D. Ashley, *Artificial Intelligence and Legal Analytics. New Tools for Legal Practice in the Digital Age* (Cambridge University Press 2017) 234.

61 Ashley (n 60) 236.

Once the source text is properly prepared, it is represented as a mathematical structure (for instance, a vector space) in a model. Then the model is subject to the process of training until it produces the results that satisfy the assumed criteria. Typical legal problems resolved by ML systems are classification (assigning a label to the new data) or prediction of an event or behavior (which may also be seen as a type of classification).⁶² For instance, in information retrieval, the task may be to retrieve defined relevant information (e.g., cases decided in favor of the plaintiff). In semantic classification systems, the result may consist of classifying objects (for instance, legal provisions). ML systems may also be used for exploratory purposes, for instance, to detect repeatable patterns of data not identified yet, which may indicate non-accidental regularities not identified previously (for instance, fraud or tax evasion). The performance of an ML model is assessed against a set of standard criteria such as precision (the ratio of the amount of true positive results to the sum of true positive and false positive results), recall (the ratio of the amount of true positive results to the sum of true positive and false negative results), traditional F-measure (harmonic mean of precision and recall), and other criteria.

Computational intelligence systems may generate erroneous results by nature, especially if the target dataset differs in certain respects from the training set. Conversely, increasingly often, the performance of ML tools for certain tasks is comparable to, or even exceeds, the level of human lawyers regarding accuracy. In particular, one may enumerate the experiment concerning reviewing contracts,⁶³ applications of the question answering system to provide legal texts relevant for legal queries,⁶⁴ pre-

62 It should be stressed that both classification and prediction tasks may also be performed by symbolic models, in particular by case-based reasoning systems and argumentation systems. See, for instance, Kevin D. Ashley and Stefanie Brünighaus 'Automatically Classifying Case Texts and Predicting Outcomes' (2009) 17(2) *Artificial Intelligence and Law*: 125–65 and the dissertation of Matthias Grabmair, *Modeling Purposive Legal Argumentation and Case Outcome Prediction Using Argument Schemes in the Value Judgment Formalism* (University of Pittsburgh 2016) <<http://d-scholarship.pitt.edu/27608/>>, accessed 17 August 2021. Nonetheless, in practical applications, the computational intelligence approach is prevalent, because of the possibilities concerning application of the model to the new datasets expressed in natural language.

63 See <<https://www.lawgeex.com/>> accessed 10 May 2021.

64 See <<https://www.rossintelligence.com>> accessed 10 May 2021.

dictions of European Court of Human Rights decisions,⁶⁵ or recent predictions concerning domain dispute decisions in the legal framework of WIPO.⁶⁶ The abovementioned systems are targeted to perform strictly defined tasks typically restricted to particular domains, but the constant evolution of the ML and NLP technologies creates possibilities for generalizations. In particular, the results obtained in the initial stage of the Lex Rosetta project show that similar, promising results may be obtained in the performance of tasks concerning the segmentation of judicial opinion issues in jurisdictions and drafted in languages.⁶⁷

However, the efficient operation of computational intelligence ML systems in the performance of legal classification and prediction tasks does not mean that their results are readily applicable to solving such legal problems as justifying an opinion, establishing the relative weight of arguments, or explaining why a situation should be regarded as an instance of an abstract concept. Even if the results generated by the numerical model are likely to be evaluated as correct by most professional lawyers, this does not mean that they were obtained along the same line of reasoning that a human lawyer or an idealized Hercules judge would present. The contrary is the case, as the operation of such systems is typically based on dozens, hundreds, or thousands of features captured by a numeric model. Nowadays, one of the most intensively debated topics in AI is its explainability: the possibility of presenting the mechanism of algorithm operations in a manner understandable to humans.⁶⁸ As we have noticed above, the symbolic AI models of legal reasoning realize this postulate to a high degree. This does not hold, however, for computational intelligence systems, the level of explainability of which varies across models and is the lowest regarding multi-layered artificial neural networks. The relatively low level of their explainability means that it is difficult, in some cases practically impossible, to answer why the system generated a given result.

65 Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, Vasileios Lamps, 'Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective' (2016) 2 *Perrj Computer Science*, e93.

66 L. Karl Branting, Craig Pfeifer, Bradford Brown, Lisa Ferro, John Aberdeen, Brandy Weiss, Mark Pfaff, Bill Liao, 'Scalable and explainable legal prediction (2020) Artificial Intelligence and Law', <<https://doi.org/10.1007/s10506-020-09273-1>> accessed 10 May 2021.

67 Jaromír Šavelka, Hannes Westermann and others, 'Lex Rosetta: Transfer of Predictive Models across Languages, Jurisdictions and Legal Domains' (ICAIL 2021: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law, São Paulo, 21-25 June 2021) 129-138.

68 See Arrieta and others (n 25).

The problem of the explainability of ML systems gives rise not only to epistemic problems but also to ethical and legal ones. If a system is used as an element of the decision-making process, we should be able to provide transparent reasons for the adoption of such a decision. The lack of such transparency and accountability may lead to (the risk of) legal liability.

These problems led to the emergence of a subdomain of XAI (explainable artificial intelligence) research, namely, the concept of XAILA (explainable artificial intelligence and law).⁶⁹ One of the main ideas discussed in this field is to bridge the gap between data-driven numerical ML systems on the one hand and the knowledge-based, symbolic AI systems on the other, and possibly to combine them in hybrid systems.⁷⁰ Such systems should aim to balance the performance of computational intelligence systems with the relatively high level of explainability of symbolic models of legal reasoning. One of the approaches represented in this field is to enhance the explainability of ML models by training them based on the data annotated with categories characteristic of the knowledge elements employed in the computational models of legal reasoning, such as legal norms, concepts, premises of arguments, inference links, etc. Such systems could explain their classifications and predictions through the recourse of the features specified in the annotation, which correspond to the intelligible elements of legal reasoning.⁷¹ The concept of combining the ML approach and the symbolic models of reasoning approach has been elaborated at a deep level in the conception of cognitive computing legal apps (CCLA) advocated by Kevin D. Ashley.⁷² The CCLA consists of forming legal hypotheses (such as a given set of circumstances that should or can lead to a given result) and then testing them in the environment encompassing the ML model, the computational model of legal reasoning, and the human. According to this approach, legal datasets should be annotated with schemes determined by the structure of computational models of legal reasoning. Therefore, they could serve as the source of premises for the latter models, which would then perform highly reliable reasoning based on valid or at least well-defined inference patterns. It is emphasized that the presence of a human in the loop is essential here, particularly because the set of premises retrieved by the ML models may be imperfect

69 A series of workshops attached to the JURIX International Conference on Legal Knowledge and Information Systems, in 2018, 2019, and 2020. See <<https://www.geist.re/xaila:start>> accessed 10 May 2021.

70 See Atkinson, Bench-Capon, Bollegala, (n 58).

71 See L. Karl Branting and others (n 66).

72 See the extensive elaboration of the idea of CCLA in Ashley (n 285) 350-391.

for various reasons. The reasoning performed by computational models of reasoning may also require verification. The availability of the CCLA could substantially enhance the performance of legal practitioners through the facilitation of data analysis (via the ML component) and ensuring correct reasoning (via the computational model of reasoning). Nonetheless, the tension between the limited rationality of human reasoners and the tendency of computational models of reasoning to rationalize them has its bearing on the ML-based prediction of legal decisions and the CCLA concept. Should we predict an imperfectly reasoned (even erroneous) human decision or the decision of an entity exceeding humans, regarding intellectual capabilities? Moreover, can such capabilities of the human mind as reasonable judgment be well defined in the sense of problem-solving theory?

These questions lead us to the problem of the fundamental dichotomy of ML models on the one hand and the models representing symbolic reasoning and justification on the other. The essence of ML models is that they represent the structure of existing data. Nevertheless, the substantial feature of legal reasoning is its normativity, understood here as the possibility of subjecting any legal claim to critical scrutiny. Irrespective of the existing practice (documented by the available sources), lawyers have the vocation to challenge it by asking whether this practice should be continued. In fact, arguments based on established practices or customs are only one type of argument among many used in legal discourse, and there is an ineliminable tension in legal reasoning between the value of stability and certainty, on the one hand, and flexibility and evolution, on the other hand. These dynamic tendencies may also be recorded in the available data. However, lawyers may also critically assess the character of these dynamics. In addition, in the Anglo-American legal culture where the evolution of case law is constrained by the *stare decisis* principle, lawyers may add dynamics to the evolution of the legal domain through creative distinguishing argumentation or, in certain situations, through arguing for overruling of earlier precedents. In statutory interpretation, this tension is captured by the potential conflict between linguistic arguments and purposive (teleological) arguments. The data-oriented nature of computational intelligence systems causes their inability to capture this normative, or open, character of legal argumentation. As this is a natural feature of these systems, it should not serve as the basis for their critique; it is simply not fit for the purpose of modeling normative aspects of legal argumentation. Contrarily, symbolic AI models of legal reasoning may present a line of argument similar or indirectly translatable to the line of reasoning that could be presented by lawyers in natural language, including the mechan-

isms of the construction of new arguments from the database. Of course, the relevant information must already be stored in the database, and the patterns of inference must be captured by the mechanisms implemented in the program. The limited or lack of ability of symbolic AI systems to generalize beyond available knowledge should not be the basis of critique of these systems. They are simply not fit for this purpose. However, they are designed to represent reliable, valid reasoning from a well-structured set of premises.

Therefore, the CCLA concept aims to draw benefits from the strong sides of both components (ML models and symbolic AI reasoning systems) and simultaneously relate the training process of the former to the elements considered relevant for the latter. Considering the radically different character of both components, the conception is a far-reaching attempt to align their operation. The presence of a human on the loop is an indispensable element of this conception because it is necessary to critically evaluate the input to the reasoning system provided by an ML model and to investigate whether the reasoning performed by the computational, symbolic system does not lead to oversimplifications. The output generated by the CCLA, concerning, for instance, predictions of outcomes, assessments of the strength of arguments, or indications of the relevant existing case law may and should be evaluated by human lawyers who may continue the iterative process by modifying questions posed to the system or proposed hypotheses submitted for verification. Undoubtedly, the development of any CCLA is a complex task, beginning with the preparation of an annotation scheme based on elements relevant to symbolic models of legal reasoning.

Another approach to the design of hybrid applications combining symbolic reasoning and ML-based task solutions can be outlined as follows. Generally, for any legal problem, possible answers may be deliberated, and justifications supporting these answers may be constructed. These alternative justifications could be generated automatically from the database encompassing both general jurisprudential knowledge (types of legal norms and legal concepts, interpretive canons, patterns of inference, catalogues of legally relevant values) and domain-specific knowledge. Then, the alternatives could be tested regarding their resistance to attacking arguments. Such a testing process may, in principle, be realized by the reinforcement learning algorithm, where agents compete to produce the best possible justification of a given legal solution. Then, this solution may be compared to the solution predicted by an ML mechanism trained on textual data. This type of application would also require humans on the loop to critically evaluate the generated justifications. It would also require the develop-

ment of a cross-domain corpus of legal knowledge, a task that was initiated decades ago but still requires extensive international and interdisciplinary collaboration.⁷³

4. Conclusions

The LegalTech community should become aware of the achievements of AI and law research, including the identified limitations of approaches and the obstacles hindering the wider practical application of some prototypical systems. This direction of information flow should enable LegalTech to avoid reinventing the wheel and to increase awareness of the complexities related to legal knowledge representation, legal reasoning, and models of classification and prediction. Neither is it the case that the symbolic AI tools at our disposal match the complexity of actual legal justificatory reasoning; nor does it hold that the application of ML tools, including NLP, can always lead to reliable, replicable, practically useful, and theoretically well-founded results. Yet, the legacy of almost five decades of AI and law research provides a firm foundation for the development of new types of legal applications, including the CCLA briefly commented on above. If LegalTech embraces this legacy, it may avoid entering dead ends, concerning, for instance, knowledge acquisition bottlenecks, computational tractability problems, or undue simplifications in both knowledge engineering and developing data mining models. The complexity of legal reasoning has not been completely accounted for in AI and law research. While the theoretical foundations thereof need continuous development, LegalTech should at least become aware of the problems that already have a computational implementation, such as procedural aspects of argumentation, reasoning with burdens and standards of proof, aspects of case-based reasoning, or theory construction based on the notion of coherence.

Problems related to AI's understandability, explainability, transparency, and eventually trustworthiness pose particularly pressing problems, as apparently a major part of LegalTech solutions are based—for reasons of performance level and scalability—on ML models. The developers and users

73 In this chapter we focus on AI applications in connection with judicial decision-making and legal argument. One of the fields of AI application in the context of law, which we have not discussed here, but which is definitively worth mentioning, also due to its interdisciplinary character, is the support of dispute resolution. See John Zeleznikow, 'Using Artificial Intelligence to provide Intelligent Dispute Resolution Support' (2021) 30 Group Decision and Negotiation 789–812.

of these systems should become aware of the state of debate concerning explainable AI and law and the conceptions concerning the relationship between the explanation of the algorithms' operations and the justificatory argumentation representing the reasons for accepting a given conclusion. This debate is far from concluded, and its practical importance is enhanced by the regulatory actions undertaken by EU authorities and related debates concerning the ethics of AI use and operations.⁷⁴ The LegalTech community should also recognize problems concerning the normative and open character of legal argumentation, which remains in tension with the data-driven approach characteristic to ML models. In this connection, it is worth analyzing for LegalTech developers where and how the role of the human reasoner is placed in the new conceptions advanced in AI and law, such as the idea of the CCLA.

The above comments are not intended to imply the informational flow between AI and law and LegalTech should be unidirectional. The contrary is the case: The practical approach of the latter may provide very valuable empirical input for the former, especially on the level of identifying the actual needs of legal practice and the processes of evaluating prototypical solutions. The LegalTech sector provides a platform for large-scale experiments of the tools and solutions that may be elaborated on the basis of or already available in the results of AI and law research. Moreover, the availability of large datasets in settings relevant to LegalTech enables the development of more realistic and generalized models, both in the field of modeling legal reasoning and computational intelligence for legal classification and prediction.

I am convinced that the actual progress of LegalTech research and applications toward enhancing the performance of actual legal problem solving involves the adoption of a more comprehensive, interdisciplinary approach. As noted, although part of AI and law research is grounded in legal-theoretical work on legal reasoning, a much more intensive, bidirectional flow of information is needed between these two fields. If such communication is absent, research on AI and law focuses on the formal

74 See for instance the Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), <<https://ec.europa.eu/newsroom/dae/items/709090>> access 10 May 2021 and the earlier documents: Ethics Guidelines for Trustworthy Artificial Intelligence, <<https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>> access 10 May 2021 and Framework of ethical aspects of artificial intelligence, robotics and related technologies, <https://www.europarl.europa.eu/doceo/document/TA-9-2020-0275_EN.html> access 16 March 2021.

and computational aspects of the developed models, leaving the specificity of legal reasoning and, generally, the performance of legal tasks in the margin. This results in the development of (overly) idealized models or in the decreased understandability of models for the lawyers. A more intensive flow of information is needed both on the level of general jurisprudence (theories of legal validity, interpretation, applications of law, etc.) as well as on the level of doctrines related to domains of law. Conversely, legal theory should become more aware of the nuanced character of AI and law research, which should not be inadequately equaled with a revival of “mechanistic jurisprudence.”

However, to contribute to the more realistic computational models of legal reasoning, legal theory should become more integrated with the interdisciplinary field of studies on mind and cognition, that is, cognitive science.⁷⁵ The research area referred to as *cognitive science and the law* has attained important status in the legal–philosophical landscape, analyzing, for instance, the role of heuristics and biases in legal reasoning. Nevertheless, a significant part of the work still needs to be done, especially in the sphere of theorizing about legal reasoning in terms of mental representations and the operations performed on them. Such research may lead to a better understanding of legal concept acquisition and formation, the actual patterns of legal rule-based and case-based reasoning, as well as the relationship between the intuitive, fast system of the mind and the slow, deliberative system. The theoretical and empirical results in this field could provide feedback to both legal theory and AI and law to finally inform LegalTech about the structure of the models effectively supporting or simulating legal thinking.

75 This direction has been already elaborated in Giovanni Sartor, *Legal Reasoning*, (Springer 2005), however it definitively needs further, interdisciplinary investigations.

Figure 1. Model of information flows between the “nearest neighbor” areas.

The figure presented above indicates the “nearest neighbor” relationships between the indicated areas; it is presumed here that the bidirectional flow of information is perhaps the most natural between these pairs due to the overlap of conceptual schemes or shared aims. However, the direct flow of information is possible between each pair of these fields. The main subject of this paper is the possible influence of AI and law on LegalTech. Nonetheless, as noted, the opposing direction of impact is also possible and potentially fruitful. LegalTech is most naturally oriented toward providing results for legal practice. Cognitive science has occupied an important position as a subfield of contemporary legal theory, and as an empirically oriented research area, it also concerns realities of legal practice, especially through psychological investigations. AI and law research has been partially rooted in the contributions of legal theory (and domain doctrines), but as discussed, this mutual link should be strengthened to benefit the quality of LegalTech applications, and for the sake of development of AI and Law solutions which would serve the society best.⁷⁶

76 See Bart Verheij, ‘Artificial intelligence as law. Presidential address to the seventeenth international conference on artificial intelligence and law’ (2020) 28 *Artificial Intelligence and Law* 181-206.