Artificial Intelligence and Law (2020) 28:267–280
https://doi.org/10.1007/s10506-020-09267-z

CONFERENCE REPORT

ICAIL Doctoral Consortium, Montreal 2019

Michał Araszkiewicz[1] · Ilaria Angela Amantea[2,3,4] · Saurabh Chakravarty[5] · Robert van Doesburg[6] · Maria Dymitruk[7] · Marie Garin[8] · Leilani Gilpin[9] · Daphne Odekerken[10] · Seyedeh Sajedeh Salehi[11]

Published online: 26 May 2020

Abstract
This is a report on the Doctoral Consortium co-located with the 17th International Conference on Artificial Intelligence and Law in Montreal.

Keywords ICAIL 2019 · Doctoral Consortium · Conference report

1 Introduction

The 17th International Conference on Artificial Intelligence and Law took place in Montreal, Québec, Canada, between 17 and 21 June 2019. The main conference was accompanied by 11 events such as workshops and tutorials. One of these accompanying events was the Doctoral Consortium, which was organized at the ICAIL conference for the third time, following the successful editions co-located with the 15th ICAIL in San Diego and the 16th ICAIL in London. The

✉ Michał Araszkiewicz
michal.araszkiewicz@uj.edu.pl

[1] Uniwersytet Jagielloński, Kraków, Poland

[2] University of Torino, Turin, Italy

[3] University of Bologna, Bologna, Italy

[4] University of Luxembourg, Esch-sur-Alzette, Luxembourg

[5] Virginia Tech, Blacksburg, VA, USA

[6] University of Amsterdam, Amsterdam, The Netherlands

[7] Faculty of Law, Administration and Economics, University of Wrocław, Wrocław, Poland

[8] Centre de mathématiques et de leurs applications, CNRS, ENS Paris-Saclay, Université Paris-Saclay, Saint-Aubin, France

[9] MIT, Cambridge, MA, USA

[10] Utrecht University and Dutch National Police, Utrecht, The Netherlands

[11] Vrije Universiteit Brussels, Brussels, Belgium

2019 edition of the Doctoral Consortium was organized by Michał Araszkiewicz from the Jagiellonian University in Krakow. The session was held on 17 June 2019 and it was chaired in Montreal by Enrico Francesconi, because unfortunately Michał was not able to attend in person, but only by video link.

The main aim of the Doctoral Consortium is to promote the exchange of ideas from Ph.D. researchers in the area of Artificial Intelligence and Law, and to provide them with an opportunity to interact and receive feedback from leading scholars and experts in the field. For this edition of the Doctoral Consortium, eight papers were submitted and, upon being reviewed by the Program Committee, all of them were accepted for the presentation at the event. All the participants have contributed a section on their research to this paper.

The remainder of this paper contains brief presentations of research communicated at the Doctoral Consortium. This overview of the projects enables the reader to appreciate the diversity of topics that feature in the AI and Law community as the topics for Ph.D. programmes. It also gives an insight into which subjects and methods are currently perceived to be important for this domain of research and thus provides the basis for reflection on the future directions of the discipline. The subjects range from computational modelling of potestative concepts through the problems of explainability of automated judicial proceedings, developing an argumentative inquiry agent in the law enforcement domain, to the design of tools automatically summarizing legal depositions. Also several topics related to broader issues concerning the social (ethical, legal) context of intelligent systems operation, such as dispute resolution, autonomous machines, privacy protection and risk management in the healthcare sector are attracting the attention of aspiring scholars. One of the most important pressing questions for the AI and Law community is how computational models of legal reasoning may be fruitfully used to tackle these socially important subjects.

Taking into account its performance in the review process and the overall quality, the paper written by Robert van Doesburg, entitled *The False, the Former and the Parish Priest* was awarded the Best Doctoral Consortium Paper prize.

We are grateful to the people who helped to organize the Doctoral Consortium or otherwise contributed to make it a successful event: the ICAIL 2019 Organization Chair, Karim Benyekhlef, the Program Chair, Floris Bex, all the members of the Doctoral Consortium Program Committee, the experts involved in the best paper selection and the participants to the session. Last but not least we are grateful to the authors of the submitted papers and their scientific advisors.

The following sections are each authored by a student participant in the Doctoral Consortium.

2 Methods and tools for analysis and management of risks, compliance and privacy in the healthcare sector: hospital at home: *Ilaria Angela Amantea*

My research work focuses on business process modeling in a service of hospitalization at home (HaH) of patients with severe disease. In particular, the research questions investigate how to balance the organization and reorganization in terms of time, resources and costs with the necessary risk management in the sense of procedural risks and regulatory compliance analysis in ICT technologies applied in the management of care.

The functioning of the service relates to separate models, managed by different persons in different offices. But each of them depends and impacts on the others. Moreover, the HaH is a new type of hospitalization that differs from traditional services in terms of the problems relating to organization, structure, cost, risk, safety, transparency and regulations that it needs to address (Aimonino Ricauda et al. 2010; Rainero et al. 2017).

Furthermore, the complexity of the problems posed by telemedicine, teleassistance and more generally by Artificial Intelligence in healthcare requires a new set of methodologies to meet both organizational management issues (i.e. concerning resources, costs and time) and normative and regulatory aspects, such as risks, attribution of responsibility, transparency, information security, and data privacy (Chartier 2014; Fishman 2013; Rose et al. 2008; Vincent et al. 2000).

My work starts from the analysis of real processes improved by technological innovations (e.g., teleassistance, telemedicine, chatbot) from several industrial partners involved in an EU funded project of the Piedmont Region in Italy, CANP1. This allows the use of real data from the process, as well as enabling a validation of the proposed business analysis from hospital staff.

Adopting a process-centric point of view, the first goal is to create the As-Is Model of the different processes. Once the processes are specified in the standard language BPMN.[1] the following step is to produce a simulation to address *what-if* scenario analysis. We plan to simulate the main HaH processes by using one of the most used discrete-event simulation tools available in the Computer Science Department of the University of Torino: iGrafx Process2. Moreover, the output of the simulation can be further investigated to address tasks of the process for monitoring procedural errors, by counting errors during the business process execution (Amantea et al. 2018, 2019).

A second purpose is to take into consideration also the legislative aspects and develop methods for representing laws and regulations directly in the specification of processes in the health sector [Risk, Compliance and Privacy by Design, Racz et al. (2010), Vincent (2011)]. We consider the laws existing today, but take into account that the law can change. The legal context is always evolving, so continuous monitoring and compliance is especially important for risks in healthcare.

[1] Business Process Model And Notation. https://www.omg.org/spec/BPMN/2.0/ accessed 20.02.2020.

The last steps in my research will further investigate both process modelling and legislative aspects, trying to provide a holistic view of how to develop an innovative service strategy which will have a positive impact on the experience of the different actors.

3 Summarizing legal depositions: *Saurabh Chakravarty*

Legal depositions are an important procedural mechanism used in civil lawsuits. Lawyers use them to discover and develop the information known by parties and other witnesses about the facts pertaining to a case. The transcripts (records) of those legal depositions are then used by others, including other lawyers and paralegals, to ascertain from them the facts pertaining to a case. These documents capture the conversation between the lawyer and the witness, which is in the form of questions and answers given under oath. Having an effective AI-based system to summarize legal depositions, producing consistent and high quality summaries, would help the legal community, which often must quickly process and disseminate the key information in a deposition. This would allow lawyers and paralegals to save time and expense, and focus on other aspects of their work.

Applying current automatic summarization methods to these documents results in low quality summaries. Due to the lack of a large amount of training data, it is a challenge to use deep learning based summarization methods to train a system in an end-to-end fashion. We propose a solution in the form of a pipeline of components, each addressing a sub-problem. We have developed a series of components, starting with one to extract the texts from deposition files with various formats. Another component handles anonymization, since depositions may have confidential and personally identifying information, including about health matters. Another key component that was developed was the dialog act classification (Chakravarty et al. 2019) of question-answer pairs, according to the ontology we devised for legal depositions, so that tailored summarization methods can be applied to each of the types of dialog act that are found in a deposition. These dialog acts can also be used to transform a question-answer pair to a simple or canonical form that can be used further in other downstream tasks to generate summaries.

From an evaluation perspective, we plan to use the ROUGE metrics (Lin and Och 2004) for the generated summaries along with a subjective evaluation by human actors for readability, coherence, coverage, and non-redundancy. A long deposition could lead to many different equally good summaries, depending on linguistic style, and on the issues of interest in the rest of the case. Even if a summary is available for a deposition to compare with an automatically generated summary, quantitative techniques now in use are poor predictors of the value of a summary; however, qualitative assessments by experts are time consuming and very expensive. Our approach will improve continuously—by enlisting the aid of personnel trained in the law in formative evaluation of the results of each step in our pipeline, as well as of generated summaries—if we use feedback both as input to deep learning and as guidance to enhance our methods.

The results of our work will include a corpus of legal depositions, along with generated summaries and related evaluation assessments, as well as an ontology of dialog acts, classifiers for the dialog acts, and other software that comprises the tailored summarization pipeline.

4 Explicit interpretations of sources of law using Calculemus-FLINT: *Robert van Doesburg*

Working in the domain of public services, we noticed the lack of well-structured methods for finding sources of norms relevant in a specific context, interpreting those sources of norms, and creating instructions for humans or machines based on these interpretations.

The goal of my Ph.D. project is to create a method for the interpretation of sources of norms, resulting in specifications for normative multiagent systems that can be used by humans and machines. To achieve this goal we study normative theory and methods for modeling normative systems. We developed a step-by-step plan for solving normative questions (Calculemus), and a language to make explicit interpretations of sources of norms (FLINT).

The Calculemus approach to normative questions is a simple step-by-step plan for solving normative questions:

1. Express a normative question.
2. Collect sources of norms relevant for answering the question.
3. Express your interpretation of sources of norms in a representation that can be discussed with all people that have an interest related to the question.
4. Apply interpretations in order to answer the question.
5. Compare your answer with those of others, and make a structured assessment of disputes.

These steps can be used in combination with any method for expressing the meaning of sources of norms. The FLINT representation of norms is based upon the need of organizations to give employees instructions as a basis for their actions, and as a reference for accountability. The choice of a model that expresses norms as either the power to act, or the duty to act in the future is based on the work on the fundamental concepts of legal norms by Salmond (1907), Hohfeld (1913), Kocourek (1930) and Hart (1961). In order to make representations of social systems I use the concept of institutional facts as advocated by Searle (2010).

In the FLINT language norms are expressed as:

1. Acts (consisting of an action, actor, object, precondition and postcondition),
2. Duties (consisting of a duty holder, claimant, creating act, termination act, and enforcing act), and
3. Facts (models that can be used to elaborate the circumstances in which elements of a precondition are true).

The Calculemus-FLINT method is tested by applying it to archetypical cases in the field of AI and Law (van Doesburg and van Engers 2019b), and real-life study cases in the domain of public services (van Doesburg and van Engers 2015, 2018, 2019a).

5 *Explainable artificial intelligence* in automated judicial proceedings: *Maria Dymitruk*

One of the goals of the AI and Law community is to create AI systems which are capable of being either a decision-maker in judicial proceedings or a supporting tool for human judges in deciding in legal cases.[2] AI research has undergone a long evolution, from completely explainable yet inflexible expert systems, to machine learning (ML) models, such as deep neural networks, which are effective but virtually impossible to verify (Adadi and Berrada 2018). The research on ML has thus far focused on prediction tasks but rarely on providing explanation for them.

The research on XAI (eXplainable Artificial Intelligence) plays a key role in paving the way towards successful application of AI in the justice system. Trevor Bench-Capon reasonably sums up the problem: "It has always been argued by AI and Law practitioners that the decision is of secondary importance, and what matters is the explanation (...)" (Bench-Capon 2018). When applied to the AI and Law field, XAI can be useful both in legal evidence-gathering (and its evaluation) (e.g. Verheij et al. 2016; Vlek et al. 2016 and Timmer et al. 2017) and in decision-making. The emphasis of my work is on decision-making.

In legal theory, the need for explanation has long preceded the debate on the ability of AI systems to conduct legal reasoning. In fact, human decisions admit post-hoc interpretability despite the "black box" nature of human brains. Similarly, despite the "black box" nature of contemporary ML models, they should be considered acceptable as judicial decision-makers or judge-supporting tools, but only if they are equipped with the ability to explain their decisions.

Part of the difficulty in the research on XAI lies in understanding what an explanation should actually contain. For automated judicial proceedings, the most significant is the type of explanation which makes it possible to control the reliability of the decision.[3] As a result, an AI decision should include an explanation that can be verified by a human.[4] Psychologists argue that because AI systems are regarded by their users as intentional agents, people expect the AI system to generate explanations within the conceptual and linguistic framework of human behavior explanation (De Graaf and Malle 2017). Ideally, an AI system should be able to create an explanation akin to that of a judge explaining how he/she arrived at a certain decision or

[2] Both models are referred to as *automated judicial proceedings*.

[3] It is worth mentioning in this context that many empirical studies in various fields proved that the explanation can change the users' attitudes towards the system: it can increase their confidence and trust in the systems e.g. (Herlocker et al. 2000; Sinha and Swearingen 2002; Bilgic and Mooney 2005 and Symeonidis et al. 2009), and improve their ability to correctly assess whether a prediction is accurate (e.g. Kim et al. 2016; Gkatzia et al. 2016 and Biran and McKeown 2017).

[4] I reject the possibility to use AI in order to make decisions which could not be controlled by humans.

the justification of the judicial decision.[5] Such justification should relate to the party to the judicial proceedings and make it possible to understand and accept the decision. It should also prove that during the proceedings the party has been properly heard.[6] It should be remembered that the justification ensures transparency of the court proceedings and hence public confidence in (automated) judicial institutions.

Taking the above into account, the justification provided by an AI system should include at least: (a) significant elements of factual circumstances of the case considered proved, (b) indication of the legal basis of the decision together with its analysis, and how it applies to the factual circumstances of the case, (c) reference to the arguments presented by the parties to the proceedings (if any) and explanation whether they were taken into account or not and why. As AI systems are increasingly allowed to make more autonomous decisions and we migrate greater responsibility to such systems, providing justifications of their decisions will become more and more crucial. XAI has a chance to be a remedy for the concerns about the use of AI systems in the judiciary (including algorithmic bias and the lack of transparency in public decision-making).

6 Privacy-preserving algorithms: *Marie Garin*

Among the current barriers in medical research and in the development of artificial intelligence are legal and privacy concerns regarding data sharing. The modernization of digital hospital systems will enable society to collect a large amount of epidemiological and environmental data. The latter are the necessary grounds for major changes in the traditional health system. However, several factors hinder the valorization of medico-administrative data such as the fragmentation of databases that are considered too sensitive to be shared, preventing advantage being taken of the multiplicity of data sources. It is essential that scholars from both the legal and scientific communities interact to satisfactorily define privacy in AI technologies.

This work addresses the issue of privacy in AI algorithms through the prism of differential privacy (Dwork et al. 2006) and decentralized learning (McMahan and Ramage 2017). It forms part of a transdisciplinary approach to law. Privacy is a crosscutting issue concomitantly dealt with by computer scientists, statisticians and legal professionals. Thus, various formulations of what confidentiality is proliferate and one of the major difficulties lies in bringing them together. In most privacy-preserving approaches, attempts are made to limit the risks of disclosure, such as identifying the owner of a data record. Several techniques have emerged, such as aggregation, data swapping or perturbation. Differential privacy, a measure of privacy based on information theory, is the subject of a growing consensus

[5] As a result, generic explanations and standard ML explanations [e.g. the list of "most predictive topics" used by Aletras et al. (2016)] are unlikely to be useful and acceptable. Similarly, rewriting the steps of the decision-making algorithm in natural language is not what is required.

[6] Judgment of the European Court of Human Rights of 16 November 2010 (Case of Taxquet v. Belgium).

(Chaudhuri et al. 2011; Duchi et al. 2013; Smith 2011; Abadi et al. 2016; Dwork 2008).

The first research question concerns the definition of privacy and the development of a mathematical framework to assess the risk exposure of personal data. With this we can evaluate the algorithmic potential of our models and those of the literature in order to identify the most effective architectures for ensuring privacy. Decentralized learning appears to benefit from privacy properties inherent in its architecture, and it therefore constitutes the central point of interest on this research axis (Bellet et al. 2019).

The second main open research question is the modelling of de-anonymization strategies such as database cross-checking or decision rule inversion. Characterization of these attacks would allow the development of a protocol for assessing exposure to disclosure risk. Such a protocol would provide strong guarantees for the development of privacy-preserving algorithms, particularly in the health sector where many technological advances are hampered by a lack of privacy proof of developed algorithms.

Finally, one of my thesis directors, Pierre Saurel, is a university professor in artificial intelligence and also a lawyer specialized in the law of new technologies. A development of multidisciplinary aspects between the technical and legal aspects is therefore intended. One of the objectives of this thesis is to highlight appropriate operational architectures in order to provide recommendations to lawyers and data access regulators.

7 Explanations for accountability of autonomous machines: *Leilani Gilpin*

There is an immediate need for machines to be able to explain their behavior and defend the reasonableness of their actions. This has caused an increase of technical contributions in explanatory artificial intelligence (XAI), but there has been very little work on

1. Whether these explanations are meaningful for liability and legal reasoning, and
2. Whether these explanations can help machines do tasks better.

My proposed thesis topic aims to tackle both of these questions by enforcing internal subsystem explanations. These internal subsystems will be combined together to tell a coherent story about the underlying reasons for the decisions or actions of the system. Such a story must be in a form that is understandable by other agents, including humans, and it must be able to be challenged in an adversarial proceeding. With this vision, I present a system methodology that uses internal subsystem explanations to diagnose and resolve anomalous behaviors. Internal explanations between subsystems will be used dynamically by the parts of the system to enable the detection of failure and intrusion. System-level

explanations will be also be useful to humans for engineering, legal reasoning, and forensics. My thesis hypothesis is that when the following two methodologies are incorporated:

- Subsystems that explain their behavior to their neighboring parts, and
- Committees of subsystems working on common tasks that are able to explain the behavior of their constituents; by processing explanations and deciding which premises should prevail,

then complex systems can more robustly decide which intended behavior is best, and support that claim with an explanation of why. In particular, my thesis aims to answer the following two questions:

- Can machine-generated explanations be meaningful and trusted for liability and legal reasoning?
- Can machine-generated explanations help machines do tasks better (faster, more reliably, etc).

This is important as autonomous systems start sharing control with humans (and may even take full-control of these tasks). For example, when the autonomous system takes over suddenly, the driver will ask why. When an accident happens in a car that is co-driven by a person and a machine, police officials, insurance companies, and the people who are harmed will want to know who or what is accountable for the accident. Control systems in the vehicle should be able to give an accurate unambiguous account of the events. Explanations will have to be simple enough for users to understand even when subject to cognitive distractions. At the same time, given the need for legal accountability and technical integrity these systems will have to support their basic explanations with rigorous and reliable detail. In the case of hybrid human-machine systems, we will want to know how the human and mechanical parts contributed to final results such as accidents or other unwanted behaviors.

The ability to provide coherent explanations of complex behavior is also important in the design and debugging of such systems, but it is essential to give us confidence in the competence and integrity of our autonomous counterparts.

8 Natural-language inquiry dialogue in the law-enforcement domain: *Daphne Odekerken*

A dialogue system, or conversational agent, is an agent that communicates with another agent. In my dissertation research, I study agents for inquiry dialogue, which collaborate with each other or with a human user in order to find evidence for a given claim. As an example, we are currently working on an intake agent that inquires into fraud complaints (Bex et al. 2016; Schraagen et al. 2018). The claim under discussion is that the complainant is a victim of online trade fraud. Agents for inquiry in the legal or law-enforcement domain should accurately and efficiently find the status of the claim under discussion, responding to natural language input

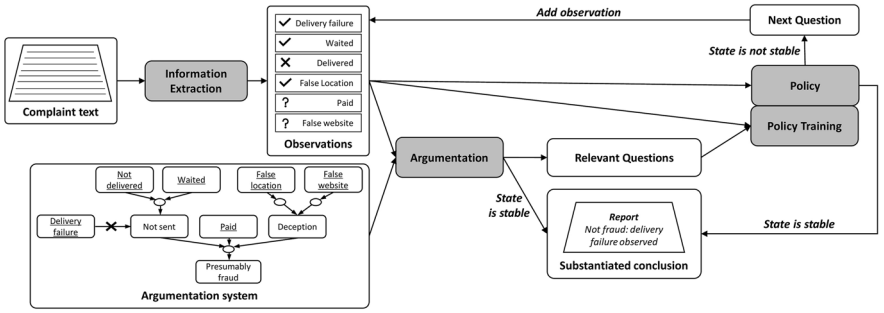

Fig. 1 Overview of the hybrid conversational inquiry agent

by asking relevant questions and drawing explainable conclusions. None of the existing methods from conversational AI (e.g. Dhingra et al. 2016; Gao et al. 2019; Jurasky and Martin 2000 and Young et al. 2013) or formal argumentation dialogue (e.g. Black and Hunter 2009; Fan and Toni 2012 and Parsons et al. 2002) meet all of these requirements. However, they have some interesting properties which can be used in a hybrid system: it is common to use machine learning techniques for handling natural language; argumentation techniques can be applied to enable transparent and accurate decision-making and asking relevant questions; reinforcement learning techniques can be used to create an efficient dialogue. In my dissertation research, I explore the possibilities of combining these approaches, answering the main research question: *How can we combine formal argumentation dialogues and machine-learning-based conversational systems into hybrid systems for inquiry dialogue and what are the advantages of these hybrid systems?*

A first hybrid agent architecture for inquiry dialogue was proposed in previous work (Schraagen et al. 2018 and Testerink et al. 2019a), and is currently being implemented at the Dutch police for the intake of fraud complaints and for cross-border information sharing. The architecture consists of information extraction, argumentation and policy learning components and is illustrated in Fig. 1.

In the **information extraction** component, machine learning techniques are applied to automatically classify observations from the initial free-text user input (handling natural language) (Schraagen and Bex 2019). The **argumentation** component reasons with these observations, based on its underlying argumentation system (Prakken 2010), which consists of a logical language and a predefined set of rules. Since we assume that the intake agent knows all the rules and the user can only contribute knowledge by stating propositions, the agent can infer which observable propositions can still change the acceptability status of the topic (Testerink et al. 2019b). In many cases, there will be several of these relevant propositions. The **policy learning** component uses reinforcement learning to find the best question to ask for any combination of observations (efficiency). When no observable proposition is relevant, the dialogue terminates. The outcome of the dialogue is accurate, because of this termination criterion and since the argumentation system is validated by domain experts. Finally, the agent is transparent since it is possible to explain for each argument supporting or attacking the topic if it should be accepted.

To conclude, hybrid agents for inquiry are interesting for further investigation. In future research, I will work on optimizing the argumentation and policy component and on comparing hybrid agents to end-to-end systems.

9 Rethinking the use of ODR in resolving e-commerce cross-border small disputes: *Seyedeh Sajedeh Salehi*

The objective of this study was to examine the role of existing online dispute resolution (ODR) mechanisms and their effects on providing consumers with convenient access to justice within the EU. The major focus of the current research was dedicated to e-commerce Business to Consumer (B2C) and Consumer to Consumer (C2C) cross-border small claims. In this regard, I elaborated the consequences of implementing ODR methods in the context of recent developments in EU regulatory safeguards to provide consumers with more effective protection against malpractices in the EU Internal Market. In this analysis, both non-judiciary and judiciary ODR redress mechanisms were considered, but the significant consideration was given to judiciary ODR methods in the EU. I specifically, investigated the implementation of the European Small Claims Procedure (ESCP) Regulation 861/2007 (as amended by Reg. 2015/2421) that was adopted with the aim of facilitating online justice for small cross-border consumer claims.

Although a considerable volume of research has already been devoted to ODR, little or no attention has been paid to providing a combined doctrinal and empirical study to assess the capability of ODR methods in resolving cross-border small claims at the EU Community level. Therefore, the methodology adopted in this research was a combination of qualitative and quantitative research methods on the basis of collecting data through conducting questionnaires and interviews within the course of carrying out the Small Claims Analysis Net (SCAN) project. As the final part of this study, I discussed the SCAN project and its significant contributions towards examining the ESCP Regulation implementation and its efficiency in providing consumers with a legal solution through using the ODR for their transnational small claims. The outcomes of this research may benefit both academia and policy-makers at national and international level.

10 Concluding remarks

We hope that the range and quality of the contributions are indicative of the health of the discipline of AI and Law. For any academic community it is vital that there is a good body of Ph.D. students coming through. The eight contributors in 2019 compares favourably with the five of 2017. But perhaps the most noteworthy feature of the papers reported here, taken as a whole, is their practical focus. All of the projects are addressing real problems that are currently being experienced by a variety of real organisations, and several of the projects are being carried out in close cooperation with the organisations that have these needs. That these projects are being driven by

perceived needs for AI and Law research rather than purely academic driven concerns surely speaks well for the future of AI and Law.

References

Abadi M, Chu A, Goodfellow I, McMahan HB, Mironov I, Talwar K, Zhang L (2016) Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pp 308–318

Adadi A, Berrada M (2018) Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). IEEE Access 6:52138–52160

Aimonino Ricauda N, Isaia G, Tibaldi V, Bestente G, Frisiello A, Sciarappa A, Cavallo S, Ghezzi M, Larini G (2011) Telecare and telemedicine in home care practice: field trial results. In: Distributed diagnosis and home healthcare, vol 2. American Scientific Publishers (ASP), Los Angeles, pp 281–303

Aletras N, Tsarapatsanis D, Preoţiuc-Pietro D, Lampos V (2016) Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective. PeerJ Comput Sci 2:e93

Amantea IA, Di Leva A, Sulis E (2018) A simulation-driven approach in risk-aware business process management: a case study in healthcare. In: Proceedings of 8th international conference on simulation and modeling methodologies, technologies and applications—volume 1 (SIMULTECH), pp 98–105

Amantea IA, Di Leva A, Sulis E (2019) Risk-aware business process management: a case study in healthcare. In: Kunreuther H, Meyer RJ, Michel-Kerjan EO (eds) The future of risk management, vol I. Springer, Berlin, pp 157–174

Bellet A, Guerraoui R, Hendrikx H (2019) Who started this rumor? Quantifying the natural differential privacy guarantees of gossip protocols. arXiv preprint arXiv:1902.07138

Bench-Capon T (2018) Legal cases: argumentation versus ML. In: ArgSoc workshop at comma 2018

Bex F, Peters J, Testerink B (2016) AI for online criminal complaints: from natural dialogues to structured scenarios. In: Artificial intelligence for justice workshop (ECAI 2016), pp 22–29

Bilgic M, Mooney RJ (2005) Explaining recommendations: satisfaction vs. promotion. In: Beyond personalization workshop, IUI, vol 5, pp 153–160

Biran O, McKeown KR (2017) Human-centric justification of machine learning predictions. IJCAI 2017:1461–1467

Black E, Hunter A (2009) An inquiry dialogue system. Auton Agents Multi Agent Syste 19(2):173–209

Chakravarty S, Phanindra RVS, Fox EA (2019) Dialog acts classification for question-answer corpora. In: Third workshop on automated semantic analysis of information in legal text (ASAIL 2019)

Chartier Y (2014) Safe management of wastes from health-care activities. World Health Organization, Geneva

Chaudhuri K, Monteleoni C, Sarwate AD (2011) Differentially private empirical risk minimization. J Mach Learn Res 12(Mar):1069–1109

De Graaf MM, Malle BF (2017) How people explain action (and autonomous intelligent systems should too). In: 2017 AAAI fall symposium on artificial intelligence for human-robot interaction

Dhingra B, Li L, Li X, Gao J, Chen Y-N, Ahmed F, Deng L (2016) Towards end-to-end reinforcement learning of dialogue agents for information access. arXiv preprint arXiv:1609.00777

Duchi JC, Jordan MI, Wainwright MJ (2013) Local privacy and statistical minimax rates. In: 2013 IEEE 54th annual symposium on foundations of computer science. IEEE, pp 429–438

Dwork C (2008) Differential privacy: a survey of results. In: International conference on theory and applications of models of computation. Springer, pp 1–19

Dwork C, McSherry F, Nissim K, Smith A (2006) Calibrating noise to sensitivity in private data analysis. In: Theory of cryptography conference. Springer, pp 265–284

Fan X, Toni F (2012) Agent strategies for ABA-based information-seeking and inquiry dialogues. Proc ECAI 2012:324–329

Fishman GS (2013) Discrete-event simulation: modeling, programming, and analysis. Springer, Berlin

Gao J, Galley M, Li L et al (2019) Neural approaches to conversational AI. Found Trends Inf Retr 13(2–3):127–298

Gkatzia D, Lemon O, Rieser V (2016) Natural language generation enhances human decision-making with uncertain information. arXiv preprint arXiv:1606.03254
Hart HLA (1961) The concept of law. Oxford Clarendon Press, Oxford
Herlocker JL, Konstan JA, Riedl J (2000) Explaining collaborative filtering recommendations. In: Proceedings of the 2000 ACM conference on Computer supported cooperative work, pp 241–250
Hohfeld WN (1913) Some fundamental legal conceptions as applied in judicial reasoning. Yale Law J 23:16
Jurasky D, Martin JH (2000) Speech and language processing: an introduction to natural language processing. Computational linguistics and speech recognition. Prentice Hall, New Jersey
Kim B, Khanna R, Koyejo OO (2016) Examples are not enough, learn to criticize! criticism for interpretability. In: Advances in neural information processing systems, pp 2280–2288
Kocourek A (1930) An introduction to the science of law. Little, Brown and Company, Boston
Lin C-Y, Och F (2004) Looking for a few good metrics: rouge and its evaluation. In: Ntcir workshop
McMahan B, Ramage D (2017) Federated learning: collaborative machine learning without centralized training data. Google Research Blog, 3
Parsons S, Wooldridge M, Amgoud L (2002) An analysis of formal inter-agent dialogues. In: Proceedings of the first international joint conference on autonomous agents and multiagent systems, pp 394–401
Prakken H (2010) An abstract framework for argumentation with structured arguments. Argum Comput 1(2):93–124
Racz N, Weippl E, Seufert A (2010) A process model for integrated it governance, risk, and compliance management. In: Proceedings of the ninth baltic conference on databases and information systems (DB&IS 2010). Citeseer, pp 155–170
Rainero C, Secinaro S, Nave E, Bignamini E (2017) Home tele-monitoring: economic and social impact of the service for patients with chronic respiratory deseas. In: 7th global innovation and knowledge accademy (Gika). Editorial ARANZADI, pp 441–442
Rose GA, Khaw K-T, Marmot M (2008) Rose's strategy of preventive medicine: the complete original text. Oxford University Press, Oxford
Salmond JW (1907) Jurisprudence: or the theory of the law. Stevens and Haynes, London
Schraagen M, Bex F (2019) Extraction of semantic relations in noisy user-generated law enforcement data. In: 2019 IEEE 13th international conference on semantic computing (ICSC). IEEE, pp 79–86
Schraagen M, Testerink B, Odekerken D, Bex F (2018) Argumentation-driven information extraction for online crime reports. In: International workshop on legal data analysis and mining (LeDAM 2018): CEUR workshop proceedings
Searle J (2010) Making the social world: the structure of human civilization. Oxford University Press, Oxford
Sinha R, Swearingen K (2002) The role of transparency in recommender systems. In: CHI'02 extended abstracts on human factors in computing systems, pp 830–831
Smith A (2011) Privacy-preserving statistical estimation with optimal convergence rates. In: Proceedings of the forty-third annual ACM symposium on theory of computing, pp 813–822
Symeonidis P, Nanopoulos A, Manolopoulos Y (2009) Moviexplain: a recommender system with explanations. In: Proceedings of the third ACM conference on recommender systems, pp 317–320
Testerink B, Odekerken D, Bex F (2019a) AI-assisted message processing for the Netherlands national police. In: ICAIL 2019 workshop on AI and the administrative state (AIAS 2019). CEUR
Testerink B, Odekerken D, Bex F (2019b) A method for efficient argument-based inquiry. In: International conference on flexible query answering systems. Springer, pp 114–125
Timmer ST, Meyer J-JC, Prakken H, Renooij S, Verheij B (2017) A two-phase method for extracting explanatory arguments from Bayesian networks. Int J Approx Reason 80:475–494
van Doesburg R, van Engers T (2015) Arguments on the interpretation of sources of law. In: AI approaches to the complexity of legal systems. Springer, pp 478–492
van Doesburg R, Van Engers T (2018) Using formal interpretations of legal sources for comparing the application of exclusion clauses of the un refugee convention. Jusletter IT, pp 175–84
van Doesburg R, van Engers T (2019a) Explicit interpretation of the dutch aliens act. In: Proceedings of the workshop on artificial intelligence and the administrative state
van Doesburg R, van Engers T (2019b) The false, the former, and the parish priest. In: Proceedings of the seventeenth international conference on artificial intelligence and law, pp 194–198
Verheij B, Bex F, Timmer ST, Vlek CS, Meyer J-JC, Renooij S, Prakken H (2016) Arguments, scenarios and probabilities: connections between three normative frameworks for evidential reasoning. Law Probab Risk 15(1):35–70

Vincent C (2011) Patient safety. Wiley, New York
Vincent C, Taylor-Adams S, Chapman EJ, Hewett D, Prior S, Strange P, Tizzard A (2000) How to investigate and analyse clinical incidents: clinical risk unit and association of litigation and risk management protocol. BMJ 320(7237):777–781
Vlek CS, Prakken H, Renooij S, Verheij B (2016) A method for explaining bayesian networks for legal evidence with scenarios. Artif Intell Law 24(3):285–324
Young S, Gašić M, Thomson B, Williams JD (2013) Pomdp-based statistical spoken dialog systems: a review. Proc IEEE 101(5):1160–1179