



AI, Law and beyond. A transdisciplinary ecosystem for the future of AI & Law

Floris J. Bex^{1,2}

Accepted: 29 April 2024
© The Author(s) 2024

Abstract

We live in exciting times for AI and Law: technical developments are moving at a breakneck pace, and at the same time, the call for more robust AI governance and regulation grows stronger. How should we as an AI & Law community navigate these dramatic developments and claims? In this Presidential Address, I present my ideas for a way forward: researching, developing and evaluating real AI systems for the legal field with researchers from AI, Law and beyond. I will demonstrate how we at the Netherlands National Police Lab AI are developing responsible AI by combining insights from different disciplines, and how this connects to the future of our field.

1 Introduction

It is my pleasure and honour to be here today in Braga¹ with many of you in the room, and many others connecting online for the first ever hybrid ICAIL conference. We live in exciting times for AI and Law. First, GPT4 passes the US bar exam (Katz et al. 2023), and not long after this a US lawyer is reprimanded for referring

¹ This text is an adapted version of the IAAIL presidential address delivered at the 19th International Conference on Artificial Intelligence and Law (ICAAIL 2023) in Braga, Portugal at the University of Minho, June 22, 2023). As a companion to this paper, I recommend the position paper presented at the 2023 Conference on Cross-disciplinary Research in Computational Law (Bex 2024), which discusses different case studies and focuses more on AI & Law as a whole instead of the work presented at ICAAIL alone.

✉ Floris J. Bex
f.j.bex@uu.nl

¹ Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

² Tilburg Institute for Law, Technology, and Society, Tilburg University, Tilburg, The Netherlands

to non-existent cases GPT has hallucinated in his GPT-generated brief.² More generally, many feel that generative language models like GPT can finally deliver on AI's promises, whereas others – including 2019 ICAIL keynote speaker Yoshua Bengio – are pointing to the dangers of such models, urging for the development of robust AI governance and regulation.³

At this ICAIL, we have seen many great presentations on completely different subjects, from trauma detection to crime simulation, to formal legal argumentation⁴. And we have heard two invited speakers with inspiring, but also cautionary tales. Natalie Byrom showed us the effects of lay people not having access to justice.⁵ Daniel Ho very vividly demonstrated what happens if the administrative state clogs up, with rows upon rows of documents in the cafeteria of government organisations.⁶

So many dramatic developments, illustrating both the positive and the negative sides of AI & Law. The question is, how can we, the AI & Law community, deal with this “algorithmic drama” we are experiencing? I will talk about three important aspects of a way forward for AI & Law: combining knowledge & data in AI, evaluating how AI & Law is used in practice, and combining multiple disciplines such as law, AI and beyond. In the rest of this address, I will show examples of how we have been tackling these points head-on in our research at the Netherlands National Police Lab AI, and I will discuss whether and how we as an AI & Law community have been engaging with these points, providing examples from this year's ICAIL conference.

2 35 years of AI & Law

But before looking ahead, let's briefly reflect on the past, as many presidents have done before me.⁷ ICAIL started in 1987 with a well-attended conference in Boston. During the 1990s, it was a solid conference with around 50 submissions and just under 100 participants (Fig. 1). By the year 2000, the AI winter being truly over, we saw a rise in the number of participants, but as then-president Karl Branting noted already in 2005,⁸ AI & Law had a difficult time emulating the successes of general AI. There were developments such as IBM's Watson (High 2012) and the application of deep learning to computer vision and language processing (LeCun et al. 2015), but AI and Law very much stuck to its ways – the focus was still on good-old fash-

² “A lawyer used ChatGPT to cite bogus cases. What are the ethics?” <https://www.reuters.com/legal/transactional/lawyer-used-chatgpt-cite-bogus-cases-what-are-ethics-2023-05-30/>, visited 4-7-2023.

³ “Pause Giant AI Experiments: An Open Letter”, Future of Life Institute, 22 March 2023 (<https://futureoflife.org/open-letter/pause-giant-ai-experiments/>).

⁴ See the proceedings of the conference (<https://doi.org/10.1145/3594536>).

⁵ Keynote talk “Increasing Access to Justice: The role of AI techniques”. See also (Byrom 2019).

⁶ Keynote talk “From Prototypes to Systems: The Need for Institutional Engagement for Responsible AI and Law”. See also (Lawrence et al. 2023).

⁷ See previous presidents' addresses (<http://iaail.org/?q=page/presidential-addresses-icail>), and also (Verheij 2020) and (Francesconi 2022).

⁸ See Karl Branting's presidential address “The Future of AI & Law” (http://iaail.org/sites/default/files/docs/ICAIL2005_PresidentialAddress_KarlBranting.pdf).




Fig. 1 ICAIL participants and submissions (left) and topics of accepted papers (right). Note that ICAIL 2021 was streamed online for free because of the Covid pandemic

ioned knowledge-based approaches, such as argumentation and case-based reasoning (Fig. 2).

The first significant increase in both participants and submissions was in 2017, due to a new-found hype surrounding AI and Law. This hype was partly due to the data revolution and technical developments in NLP, such as Transformers (Vaswani et al. 2017), making 2017 the first year that more papers on data-driven AI were published than papers on knowledge-driven AI (Fig. 1). “Legal prediction” became a thing (Ashley 2017), and legal tech was once again cool.⁹ Developments in data protection law such as the EU GDPR also led to renewed interest from legal practitioners and scholars. This led to an increase in papers on legal aspects of AI (Fig. 1), with 2019’s best student paper award going to such a paper for the first time (Almada 2019).

And here we are in the beautiful summer of AI & Law, at the first post-Covid, post-ChatGPT ICAIL, with a record number of 134 submissions and more than 300 participants. There are many papers on data-driven AI, and even a few on GPT (Blair-Stanek et al. 2023; Jiang and Yang 2023; Savelka 2023). We also see a stable core of publications on knowledge based-AI, and an increase of papers that combine data-driven and knowledge-based AI. And there is a further increase in papers on law and legal aspects (Hulstijn 2023; Nielsen et al. 2023; Unver 2023). So, from an academic perspective, AI & Law is doing great. In addition to the conference, our journal has also been performing well, ranking consistently in the top quartile of journals for both the field of Artificial Intelligence and the field of Law.¹⁰ There are many new related workshops and journals – like the recurring workshop on Natural Legal Language Processing (NLLP)¹¹ and the journal for Cross-Disciplinary Research in Computational Law¹² – and new connections – the Society for Empirical Legal Studies¹³ has a

⁹ See also Katie Atkinson’s 2017 presidential address “AI and Law in 2017: Turning the hype into real world solutions” (http://iaail.org/sites/default/files/docs/ICAIL2017_PresidentialAddress_KatieAtkinson.pdf).

¹⁰ According to the Scimago Journal rankings: <https://www.scimagojr.com/journalsearch.php?q=13880&tip=sid&clean=0>.

¹¹ <https://nllpw.org/>.

¹² <https://journalcrcl.org/crcl>.

¹³ <https://community.lawschool.cornell.edu/sels/>.

special track at this year's ICAIL. Furthermore, many of the community are involved in big national and international projects.¹⁴

AI & Law practice and applications – legal tech - is also on the rise. The number of Legal Tech companies in the Codex Techindex grew from 700 to more than 2000 in the last five years.¹⁵ ChatGPT is being actively used by, for example, judges¹⁶, and countries like China are investing heavily in legal AI and NLP for courts.¹⁷ Legal Tech companies are looking to further capitalise on the ChatGPT hype by developing derivative products for legal services¹⁸. And we see an increasing number of collaborations between universities on the one hand, and law enforcement, law firms and courts on the other hand.¹⁹

So, AI & Law is doing great by all accounts, right? Well, that depends. There have been serious worries about the use of predictive AI in courts, particularly about fairness, transparency and the effects on judicial authority (Stern et al. 2020).²⁰ Amnesty International is “Sensing Trouble” with respect to predictive policing in the Netherlands (Amnesty International 2020). There have been many discussions on ChatGPT, with the “out of the box” usage of the technology by lawyers being questioned.²¹ Even if the tech is improved, we are still dealing with technology that is far from transparent and we have seen a pushback against OpenAI on privacy violations.²² And finally, many prominent AI researchers – including Geoffrey Hinton²³ and the

¹⁴ For example, in the last 5 years community members have received prestigious Advanced Grants from the European Research Council (CompuLaw - <https://site.unibo.it/compulaw/en/project> - and Leds4XAIL <https://site.unibo.it/leds4xail/en>), and the Netherlands Research Council has granted two large 15 million Gravitation grants that involve several AI & Law community members (Hybrid Intelligence - <https://www.hybrid-intelligence-centre.nl/> - and AlgoSoc - <https://algosoc.org/>).

¹⁵ <https://techindex.law.stanford.edu/statistics>.

¹⁶ “Pakistani judge uses ChatGPT to make court decision”, Gulf News, 13 April 2023 (<https://gulfnews.com/world/asia/pakistan/pakistani-judge-uses-chatgpt-to-make-court-decision-1.95104528>). “Colombian judge says he used ChatGPT in ruling”, The Guardian, 3 February 2023 (<https://www.theguardian.com/technology/2023/feb/03/colombia-judge-chatgpt-ruling>).

¹⁷ “China’s court AI reaches every corner of justice system, advising judges and streamlining punishment”, South China Morning Post, 13 July 2022 (<https://www.scmp.com/news/china/science/article/3185140/chinas-court-ai-reaches-every-corner-justice-system-advising>).

¹⁸ E.g., legal contract drafting (<https://www.spellbook.legal/>) or general legal tech (<https://www.harvey.ai/>).

¹⁹ For example, multiple Dutch universities are collaborating with the Dutch Police in the National Police lab AI (<https://www.uu.nl/onderzoek/ai-labs/nationaal-politielab-ai>), Monash University and the Australian Federal Police have started an AI for Law Enforcement and Community Safety lab (<https://ailcs.org/>), and the University of Liverpool is working on legal tech with UK law firms (<https://www.liverpool.ac.uk/collaborate/our-successes/developing-ai-for-the-legal-sector/>).

²⁰ “As Malaysia tests AI court sentencing, some lawyers fear for justice”, Reuters, 12 April 2022 (<https://www.reuters.com/article/malaysia-tech-lawmaking-idUSL8N2HD3V7>).

²¹ See note 3.

²² “OpenAI’s regulatory troubles are only just beginning”, The Verge, 5 May 2023 (<https://www.theverge.com/2023/5/5/23709833/openai-chatgpt-gdpr-ai-regulation-europe-eu-italy>).

²³ “‘The Godfather of A.I.’ Leaves Google and Warns of Danger Ahead”, New York Times, 1 May 2023 (<https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>).

earlier mentioned Yoshua Bengio - have been warning about the dangers ahead if we let such generative AI systems run loose.²⁴

3 Algorithmic drama in AI & law

It seems we are in the middle of what Ziewitz (2016) calls an “algorithmic drama”. Like any good drama, there are protagonists and antagonists: autonomous Robojudge and Robocop as a threat versus supporting AI systems for judges and police officers as a friendly aid. These two sides of the issue we see more in AI. For example, we have on the one hand the deep learning adepts (e.g., LeCun et al. 2015) – all intelligent behaviour can be learnt by computers – and on the other hand those who argue that complex reasoning cannot be learnt but should be hardwired in the system (e.g., Marcus and Davis 2019). We also see two sides in the discussion on what AI will bring us. On the one hand, we have the techno-optimists who argue that AI can be used for good – it can help us in building a sustainable, fair society (e.g., Sætra 2022). On the other hand, we have the techno sceptics, who paint a picture of a “black box society”, in which data and algorithms are being used to control people and information (e.g., Pasquale 2015). And finally, there are the different academic fields focussed on, on the one hand, building and applying AI responsibly – the “AI” in AI and Law – and on the other hand, regulating and governing AI – the “Law” in AI and Law.²⁵

While a bit of drama is good, too much drama is not helpful. Even when we disagree on content, we want to work with each other, instead of against each other. Be open to other viewpoints and, more importantly, open to fundamentally different ways of looking at the world and AI’s role in it. But how can we do that?

First, by combining research on knowledge-based AI with research on data-driven AI: using new deep learning techniques, without forgetting about good old-fashioned AI & Law. Purely data-driven techniques are not always suited for legal AI, which involves making complex legal decisions in a way that is transparent, contestable and in line with the law. We have a strong tradition in (legal) computational argumentation (Prakken and Sartor 2015), case-based reasoning (Rissland et al. 2005), and semantic web technologies for the law (Casanovas et al. 2016). New community members – many of whom come with a machine learning background – have to be informed about this important and still relevant work.

Second, we must put AI & Law into practice. Developing systems for practice is hard, but it is the only way to fully evaluate the impact of the AI techniques we research. Of course, we already encourage applications via the innovative application award, and also numerous workshops focused on applications, but we see that applications are often stuck in the prototype stage without a proper user evaluation. Only by implementing systems at scale, with real users, can we figure out what their real impact is.

²⁴ See note 4.

²⁵ The AI & Law community is mostly focussed on AI, and the Law & Technology community is mostly focussed on Law.

Finally, we need work together across disciplines: bring together those who think about how to build AI and those who think about how to govern and regulate it. We have many community members who work at law schools as well those who are from AI and computer science departments. But we can go beyond AI and Law, reaching out to other disciplines such as public administration, philosophy or communication studies. With AI becoming commonplace in today's society, the application of AI to the law is no longer just of interest to one community.

4 Research and development in the national police lab AI

Let us now look at some examples of how we have incorporated the above three points in our research. Much of this research takes place in the National Police lab AI (NPAI), a collaboration between multiple universities and the Netherlands National Police.²⁶ At the lab, we research and develop AI tools and systems for real police problems, in an actual police context: of the 20 PhD candidates in the lab, the majority work part-time at the police while writing their thesis. We started in 2019 to develop AI tools for the police, so many PhD candidates have a computer science or AI background. Over the years, however, we have also involved PhD candidates from other disciplines²⁷ to broadly evaluate these AI tools: How are they used at the police? Which legal safeguards are required when the police use them?

4.1 AI for citizen complaint intake

The first example of research in the NPAI is about an AI system for citizen complaint intake (Odekerken et al. 2022). Such complaints are about online trade fraud, for example, false web shops or malicious traders on eBay not delivering products to people. The police receive about 60,000 complaint reports of alleged fraud each year, but not all of these are actual criminal fraud – someone might have, for example, accidentally received the wrong product. The problem was that the police needed to manually check all these reports. To solve this, we developed a recommender system that, given a complaint form, determines whether a case is possibly fraud, and then only recommends filing an official report if it is. This system was implemented at the police and is still in use today.²⁸

For the intake system, we wanted to combine knowledge- and data-driven AI. Knowledge-driven AI is used because the domain is bounded and the rules are known, because whether something is fraud or not depends on Article 326 of the Dutch Criminal Code and police policy rules, which can be directly modelled. We therefore made a legal model of the domain, as structured arguments in ASPIC+

²⁶ <https://www.uu.nl/onderzoek/ai-labs/nationaal-politielab-ai>.

²⁷ For example, the NPAI collaborates with the ALGOPOL Project (<https://algopol.sites.uu.nl/?lang=en>), which includes researchers from media studies and public management studies. The NPAI is also an important stakeholder in the AI4Intelligence Project (<https://www.uu.nl/en/research/ai-labs/national-police-lab-ai/projects/ai4intelligence>), which includes researchers from law and public management studies.

²⁸ <https://www.politie.nl/aangifte-of-melding-doen/aangifte-van-internetoplichting-niets-ontvangen.html>.

(Prakken 2010). Figure 2 (middle) shows a simplified example of the legal model that was implemented at the police. Here, something is fraud if the product was paid for but not sent and some form of deception was used. Two instances of deception are that the supposed seller used a false location, or a false website. We also include exceptions so, for example, if there was a delivery failure, we cannot say the product was not sent.

The complaint reporting form also contains a free text field where the citizen can tell their story, so data-driven AI is needed to extract the basic observations that act as the rule antecedents from this free text. In Fig. 2, these basic observations are underlined. We experimented with various machine learning NLP approaches to extract, for example, entities and relations from the text (Schraagen et al. 2017; Schraagen and Bex 2019). While results were acceptable, the final implementation mainly depends on regular expressions to extract observations. Using these regular expressions, we can extract basic information – in the case in Fig. 2 (left), the complainant has paid, the seller used a false location, and the product was not delivered. With this basic information, we can try and infer a conclusion. Here, this is not yet possible. Often the complainant does not give all the necessary information directly in the form, also because they do not know exactly what is relevant.

To find missing observations, the system can ask the complainant questions. The system will first try to determine which observations can still change the conclusion. In the case of Fig. 2 (right), there are three observations the system can ask for: was a false website used, was there a delivery failure, and did the complainant wait? Only the latter observation will allow us to possibly infer the conclusion “fraud”, as we already have a form of deception, namely false location, and we have no argument for “not sent” that can be undercut by the delivery failure. Determining whether the conclusion can still change, and which observations are still relevant for such a change, is computationally quite expensive, so we developed argument approximation algorithms for this (Odekerken et al. 2022).

The next important issue was how to explain the system’s conclusions to the complainant. It is often assumed that knowledge-based AI methods are “inherently explainable”, but there are many different explanations given a complex, rule-based system. In our case, we defined various types of explanations for the acceptance or

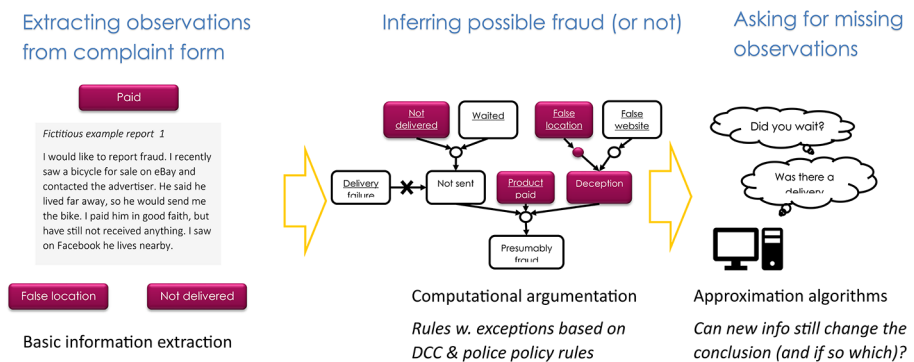


Fig. 2 Steps in the AI for citizen complaint intake process

non-acceptance of a conclusion (Borg and Bex 2020). If, for example, the complainant was to answer negatively to the questions in Fig. 2, then the system would say that the conclusion is not accepted because the “waited” observation is missing.

The intake system has been evaluated internally at the police on various aspects, such as accuracy and user satisfaction.²⁹ In our ALGOPOL project, we also wanted to further evaluate the system’s effect on human trust: would citizens mind that they received recommendations from a computer? And would it matter if they received an explanation for the recommendation or not? We performed a controlled experiment with more than 1700 participants, together with colleagues from public management studies (Nieuwenhuizen et al. 2023).

So here we had a situation where the system told the participants it was probably not criminal fraud in their case. We then asked them whether they trusted the system’s conclusion. More importantly, we also measured their trusting behaviour: the system told them it was probably not criminal fraud, but did they still file an official report? We compared two groups. The control group received no explanation – “it is probably not criminal fraud, so the system recommends you don’t file a report”. 40–60% still filed a report, so quite a large number of respondents from the control group listened to “computer says no”. With an explanation similar to the one just discussed, however, only 20–35% still filed a report, so significantly more people followed the recommendation if it was accompanied by an explanation. From this we concluded that citizen trust increases with explanations.

What this shows is how AI can be designed and evaluated from different perspectives, and how these different perspectives can inform each other. Explanations are important for citizen trust, and only this type of system can give explanations based on legal and policy rules. Here we evaluated the influence of being transparent about a single decision, taken by one AI system. This type of transparency about a single decision or a single system is what we traditionally associate with XAI – Explainable Artificial Intelligence. But it is also possible to, for example, explain how systems are being used by police officers in the organisation, or how the organisation adheres to more general regulation on data and AI. In future research, we want to test the effects of these other types of explanations and transparency about AI at the police.

4.2 Explainable AI

The second example concerns data-driven natural language processing. The police generate, use and analyse lots of text data: citizen reports like the ones from the intake system, incident reports filed by officers, forensic lab reports, intercepted communications and seized data carriers. This enormous amount of data has overwhelmed crime investigators: How can we find evidence about a certain suspect or certain event in millions of intercepted messages?

One technique for this is text classification, which can be used for searching in large document sets. For example, a classification model can indicate which out of

²⁹ The system was around 90% accurate when measured against what police case workers would recommend (submit or do not submit report). The efficiency of the reporting process was increased (from around 2 h to 0.25 h per complaint). Users (citizens) gave the system a high satisfaction rating (4+ out of 5).

the millions of messages was probably written by the suspect. But text classification can also be used in AI systems: for example, for finding observations in reports – does this report contain a mention of the fact that a product was delivered? Text classification models at the police need to be able to explain their output, as explanations can help to test and improve the model. Furthermore, new AI regulations will require the police to perform assessments of deployed models, where increasing demands on transparency require that the police understand and explain model behaviour. And if the outcome is to be used as evidence in a criminal case, explanations are also needed: “why did the algorithm mark exactly this message out of 10.000 as being very probably written by my client?”, a lawyer could argue.

One specific explanation technique that we developed recently at the NPAI concerns generating human-like rationales, or reasons, for the output of a classifier (Herrewijnen et al. 2021). In the example in Fig. 3, the model classified the top message as “paid” because it contains the sentence “I paid him in good faith”. Similarly, it classifies the bottom message as “not paid” because it contains the sentence “I haven’t transferred the money yet”. Here, the model itself has generated the rationales, and sentences instead of just individual words are highlighted, so we generate faithful rationales that make sense to humans. Another technique we have developed at the NPAI is that of generating counterfactual input text that would lead to a different classification (Roeber et al. 2021). So, for example, if the original message was classified as “paid”, the model will change the message so that it will be classified as “not paid” – in the example in Fig. 3, for example, it changes the phrase “I paid him” to “I did not pay him”, which causes the message to be classified differently. Here, we have realistic and perceptible counterfactuals.

Rationales and counterfactuals are specific explanation techniques, of which there are many for language processing, including the now-standard LIME (Ribeiro et al. 2016) and SHAP (Lundberg and Lee 2017) techniques. In AI & Law, explainable text processing has also attracted interest (e.g., Branting et al. 2021; Tan et al. 2020). With many of these specific techniques, however, the question remains exactly how they are to be used for testing, explaining and improving models?

As a first step of using such specific explanation techniques in a coherent way towards the explanation goals of model improvement, assessment and transparency, we have developed the Explabox, a collection of libraries and a toolkit for AI model

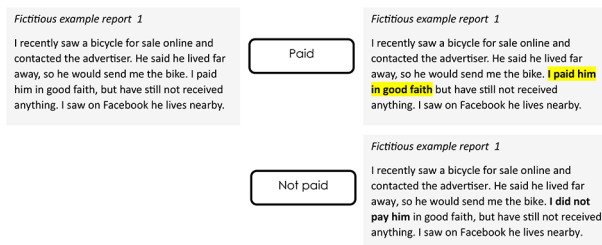


Fig. 3 Highlighted rationales that explain the “paid” classification (top), and a generated counterfactual for the opposite class “not paid”(top)

inspection (Robeer et al. 2023).³⁰ In addition to allowing for explanations via techniques such as rationales, LIME or SHAP, the Explabox allows for exploration of the data by generating basic data statistics, and for testing the robustness of a model by changing the data. For example, does the behaviour and performance of the model change when we introduce spelling mistakes, when we change all the names from Dutch to English names, or when we replace every pronoun with “she/her”? The Explabox is thus meant to give a data scientist a “holistic” view of the AI system: what kind of data was used, and what is the behaviour of the system using this and other similar data?

Issues like model behaviour, bias and fairness have become important, also in AI & Law (Alikhademi et al. 2022; Tolan et al. 2019). For some years now, there has been a special ACM conference on fairness, accountability and transparency.³¹ With respect to accountability, there is the upcoming EU AI act,³² which will – among other things - require AI in law enforcement to be certified. Regular impact assessments of AI systems will need to be performed using auditing tools like the Dutch government’s Fundamental Rights and Algorithms Impact Assessment (FRAIA)³³, which can be used to assess why and how AI is being developed and, importantly, what the impact of the AI on fundamental rights such as privacy or non-discrimination might be. The information that can be generated using Explabox can directly be used in impact assessments such as FRAIA to answer questions like “what kind of data are you using?”, “is there a bias in the data?”, “how accurate is your AI model?”, “can you explain what your model does?”.

In addition to technical research and development, we also want to further investigate the effects of the various tools and metrics to assess the impact of AI. Do they really lead to better (use of) AI? Or to a better weighing of the fundamental rights and values at stake? In the Algosoc project,³⁴ we will empirically investigate what the intended and actual effects of the current “AI audit explosion” are (cf. Power 1994). One thing we already see at the police is that new roles and responsibilities appear: instead of just the builder and user of an AI system, there is now also the auditor and the examiner of AI systems. Furthermore, we also want to look at what the law tells us about AI, transparency and contestability (Almada 2019; Bibal et al. 2021). How can legal decisions based partly on AI be motivated in a way that they are contestable and transparent? We know, for example, that argument-, scenario- and case-based reasoning are all used when motivating decisions (Atkinson et al. 2020; Bex 2011), but can AI also be explained in this way?

³⁰ <https://explabox.readthedocs.io/>.

³¹ The ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT, <https://facctconference.org/>).

³² The Proposal for a Regulation laying down harmonised rules on artificial intelligence (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>).

³³ <https://www.government.nl/documents/reports/2021/07/31/impact-assessment-fundamental-rights-and-algorithms>.

³⁴ <https://algosoc.org/>.

4.3 Evaluating AI in practice at the police

The next two examples tell us something about the importance of evaluating in a real user context, which was already indicated by Conrad and Zeleznikow (2013, 2015). The first of these extensive user evaluation projects is about a police system that is meant to detect drivers who are using their mobile phone hands-on while driving. Here, image recognition software does initial filtering for pictures of drivers who seem to be holding something that looks like a mobile phone in their hands. Whenever there is a “hit”, a human police officer checks whether the driver is really holding a phone – it might be, for example, that the phone is mounted on the dashboard or the windscreen.

During the evaluation (Fest et al. 2024) the system was found to be a best-practice in value-sensitive design. Data minimalization was considered by not storing pictures that are not tagged by the system, and only showing the driver’s side of the car. The machine learning models were trained with representative datasets, and soft- and hardware was developed in-house, so that control lies with the police and not with external commercial partners - recall that the Compass recidivism system was black box exactly because it was proprietary technology (Angwin et al. 2022).

In practice, however, there were some unexpected surprises. Many new cars had a sun-blocking windscreen foil that made it hard for the AI to detect whether someone was holding a phone. So, according to some metrics we could say the system was being “unfair” to people with older cars. Also, whenever the human officer was unsure about exactly whether the driver was holding a phone, they took a picture of the screen with their own phone and sent this to a colleague to ask for a second opinion. This was all in good faith, but it does run contrary to the data minimalization principles designed into the system at the start. So designing responsible AI, requires continuous training of the system – in the case of new windscreen foils – and the human users – in the case of assessing photographs.

A second project that led to unexpected user behaviour is AI for police interception (van Droffelaar et al. 2022). Whenever there is a crime like a robbery or a smash and grab and the suspects flee using a motorized vehicle, this system predicts the suspect’s route using knowledge about, for example, roads and suspect behaviour. This is then relayed to the dispatchers, who can tell the police cars where to best intercept the suspect. The question was whether dispatchers care about what such a system says: expert dispatchers have their own experience and knowledge about how typical suspects flee.

During the evaluation (Selten et al. 2023), it became clear that police dispatchers only followed the recommendation – the prediction – of the system when it coincided with their own intuitions. And interestingly, unlike the citizens who trusted the intake system more when they received an explanation (Sect. 4.1), the police officers were hardly influenced by explanations. So, we see that there are no “one size fits all” solutions in XAI - it is very hard to generalise over different types of users and tasks. And what we also see is that in human-AI teams, it is just as important to retrain the human as it is to retrain or redesign the AI – if, for example, this interception prediction system is more accurate or faster than humans in some cases, we would want the dispatchers to follow it.

5 AI, law and beyond

Recall that earlier (Sect. 3), I mentioned three points that are in my opinion central to a way forward for AI & Law: (1) combine research on knowledge-based AI with research on data-driven AI; (2) evaluate how AI & Law is being used in practice; and (3) combine multiple disciplines. I have shown how in the Police lab, we take these points seriously (Sect. 4). For example, in the trade fraud complaint intake system we combine knowledge-based AI and data-driven AI. We have put specific XAI-NLP-techniques into practice via a general tool like the Explabox. We evaluated different systems, including the trade fraud intake system, in practice, showing positive and negative, expected and unexpected effects of these systems. In these evaluations, we have worked with researchers from computer science, AI, law, public management, and media studies, with more upcoming projects also taking a multi- and interdisciplinary angle.

But where do we stand on each of these three points as a broader AI & Law community? Let's take a look at 2023's ICAIL conference to see whether, and how, the three points return in the accepted papers.

5.1 Combining knowledge-based and data-driven AI

Combining knowledge and data has been on the AI community's agenda for some time now (Marcus and Davis 2019; Sarker et al. 2021), and researchers in AI & law have also made first steps. First, there is the work on legal information extraction, where information that can be used in knowledge-based systems – rules, factors, arguments, cases – is extracted from unstructured data, usually text. We have seen examples of this at ICAIL 2023: Santin et al. (2023) and Zhang et al. (2023) mine argument structures from legal texts, Gray et al. (2023) automatically identify legal factors in legal cases, and finally Zin et al. (2023) and Servantez et al. (2023) extract logical formulas from legal texts. Second, there is work that uses non-statistical, knowledge-based techniques to reason with data, performing tasks such as classification, generalization or explanation (Blass and Forbus 2023; Odekerken et al. 2023; Peters et al. 2023). Finally, we have systems where data-driven and knowledge-based AI are combined in a single system. For example, when machine learning is used for information retrieval, and knowledge-based techniques are then used to reason with this information (Mumford et al. 2023) – this was also what we saw earlier in the complaint intake example for the police (Odekerken et al. 2022). Another approach is to use a combination of machine learning and knowledge-based techniques to perform one task – for example, improving court forms (Steenhuis et al. 2023). And then there are combination approaches we do not see in AI & Law or at ICAIL, yet. The first of these involves using machine learning architectures to solve typical knowledge problems, like argumentation or case-based reasoning (Craandijk and Bex 2021; Li et al. 2018). The second involves constraining what machine learning models learn, or can learn, using symbolically represented knowledge (Gan et al. 2021).

The advancements made in machine learning, and particularly in language processing and language generation, are impressive. But it can be argued that these new techniques are not always suited for legal AI that involves making complex legal

decisions in a way that is transparent, contestable and in line with the law. As argued before, we should not forget about the important work on knowledge-based AI that has been done in AI & Law. Why should we learn correlations when explicit rules and regulations already exist? And how can we guide the next generation of large language models to take existing law into account? Devising solutions to these questions is at the core of AI & Law.

5.2 Evaluating AI & law in practice

We next turn to evaluating AI & Law applications in practice. ICAIL 2023 once again has a nice amount of innovative application papers (Fuchs et al. 2023; Haim and Kesari 2023; Hillebrand et al. 2023; Steenhuis et al. 2023; Westermann and Benyekhlef 2023). However, of all of these, only Westermann et al. (2023) and Steenhuis et al. (2023) engage with humans for their evaluation, and of these two only Westermann evaluates the implemented system with actual users. So, the innovative applications are not evaluated in an operational context, on their usability or on the impact they have on legal decision making.

So, it seems that not much has changed: In 2013 and 2015, Conrad and Zeleznikow indicated that human performance or operational-usability evaluations are only presented in about 10% of papers at ICAIL and in the AI & Law journal that present some sort of application.³⁵ Work on an application together with domain users is not easy. For example, a lot of time goes into practical solutions that are not interesting for an academic publication, such as the regular expression-based information extraction in the complaint intake system (Sect. 4.1). Expert user such as law enforcement officers, judges and lawyers are busy people with little time for iterative research and development of an AI prototype system that might not be further implemented. And as we have seen in the examples in Sect. 4.3, users have their own behaviour that might give problems for even the best-designed applications.

That all being said, I firmly believe that working with stakeholders from practice is necessary in an “applied” field such as AI & Law. We cannot rightly claim we develop AI for the legal field, if ultimately only very few in that field can use our systems and techniques (or at least derivatives of them). Even if it is not possible to work with practitioners on a day-to-day basis like in the police lab, we should try to evaluate with “proxy-users” such as students. Furthermore, we can develop for user groups that are not as pressed for time as police officers and lawyers – for example, systems for legal education to students (Aleven and Ashley 1997), or for other academics such as our colleagues in empirical legal studies.

5.3 Working with different disciplines

This brings me to my final point – working across different disciplines. AI & Law is inherently multi- and interdisciplinary. Of course, we apply AI techniques to the legal

³⁵ Note, however, that an informal analysis of AI & Law journal articles from 2023 shows that at least 30% of the articles presenting some sort of application have a form of (expert) user evaluation or involvement, see (Bex 2024).

field. But we also use AI, such as agent-based simulations, to study law enforcement (van Leeuwen et al. 2023) and the law (Fratric̃ et al. 2023). At ICAIL 2023, the special Empirical Legal Studies track has led to a further increase of papers in which AI techniques are applied to legal data to study the law (Habba et al. 2023; Piccolo et al. 2023; Riera et al. 2023; Schirmer et al. 2023). And while we are not a “Law & Tech” conference that is mainly visited by legal scholars, we do have legally oriented papers at ICAIL. For example, there are articles on “law-by-design” – how legal concepts can be directly implemented in AI systems (Hulstijn 2023) – on legal aspects of our kind of AI, that is, AI for the legal sector (Unver 2023), and on the effects of AI on the legal process (Nielsen et al. 2023).

Lawyers are, in a way, still close to computer scientists. They are mathematicians with words, or social engineers. For the field of AI & Law to mature, we need to look beyond just Law and Computer Science to other disciplines. People from management or information sciences allow us to zoom out and see the bigger picture: like the socio-technical system architectures Daniel Ho talked about in his invited address at ICAIL 2023.³⁶ More critical humanities, such as philosophy but also media studies theory and science and technology studies, question some of the core behaviours and ways of communicating that we in our community take for granted (cf. the evaluation of “value-sensitive by design” in Sect. 4.3). Researchers from, for example, public administration look at our technology with an empirical lens (cf. the experiments on citizen and police expert trust in system recommendations in Sect. 4.1 and 4.3, respectively). Further strengthening ties with disciplines beyond AI and Law will ultimately benefit the entire “AI & Law” ecosystem that we are a part of.

6 Conclusion

It is the aforementioned ecosystem that I would like to conclude with. AI and Law started out in 1987 with researchers from two fields, AI or computer science, and law, studying AI for the law, using knowledge- and data driven methods. Over the years, we have involved the legal tech community and practitioners such as lawyers and educators, and started work on innovative applications and the legal aspects of our applications. The importance of broader evaluations was argued for, and more stakeholders and academic disciplines joined the field. And here we are today, with an AI & Law community – or ecosystem – that includes many more disciplines and stakeholders, and studies AI for Law and Law for AI in a broad societal framework. I am excited for the future of our field, and hope that you will help me in strengthening and further expanding this transdisciplinary ecosystem for AI & Law.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use

³⁶ See note 7.

is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aleven V, Ashley KD (1997) Teaching case-based argumentation through a model and examples: empirical evaluation of an intelligent learning environment. *Artif Intell Educ* 39:87–94
- Alikhademi K, Drobina E, Prioleau D, Richardson B, Purves D, Gilbert JE (2022) A review of predictive policing from the perspective of fairness. *Artif Intell Law*, 1–17
- Almada M (2019) Human intervention in automated decision-making: toward the construction of contestable systems. *Proc Seventeenth Int Conf Artif Intell Law* 2–11. <https://doi.org/10.1145/3322640.3326699>
- Amnesty International (2020) *We sense trouble: Automated discrimination and mass surveillance in predictive policing in the Netherlands*. <https://www.amnesty.org/en/documents/eur35/2971/2020/en/>
- Angwin J, Larson J, Mattu S, Kirchner L (2022) Machine bias. *Ethics of data and analytics*. Auerbach, pp 254–264
- Ashley KD (2017) *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge University Press. <https://doi.org/10.1017/9781316761380>
- Atkinson K, Bench-Capon T, Bollegala D (2020) Explanation in AI and law: past, present and future. *Artif Intell* 289:103387
- Bex FJ (2011) *Arguments, stories and criminal evidence: a formal hybrid theory*, vol 92. Springer Science & Business Media
- Bex FJ (2024) Transdisciplinary research as a way forward in AI & Law. *Journal of Cross-Disciplinary Research in Computational Law (CRCL)*, To appear
- Bibal A, Lognoul M, De Streele A, Frénay B (2021) Legal requirements on explainability in machine learning. *Artif Intell Law* 29:149–169
- Blair-Stanek A, Holzenberger N, Van Durme B (2023) Can GPT-3 perform statutory reasoning? *Proc Nineteenth Int Conf Artif Intell Law* 22–31. <https://doi.org/10.1145/3594536.3595163>
- Blass J, Forbus KD (2023) Analogical Reasoning, Generalization, and Rule Learning for Common Law Reasoning. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 32–41. <https://doi.org/10.1145/3594536.3595121>
- Borg A, Bex F (2020) Explaining arguments at the Dutch national police. *International Workshop on AI Approaches to the Complexity of Legal Systems*, 183–197
- Branting LK, Pfeifer C, Brown B, Ferro L, Aberdeen J, Weiss B, Pfaff M, Liao B (2021) Scalable and explainable legal prediction. *Artif Intell Law* 29(2):213–238. <https://doi.org/10.1007/s10506-020-09273-1>
- Byrom N (2019) *Developing the detail: evaluating the impact of Court Reform in England and Wales on Access to Justice*. The Legal Education Foundation
- Casanovas P, Palmirani M, Peroni S, Van Engers T, Vitali F (2016) Semantic web for the legal domain: the next step. *Semantic Web* 7(3):213–227
- Conrad JG, Zeleznikow J (2013) The significance of evaluation in AI and law: A case study re-examining ICAIL proceedings. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, 186–191. <https://doi.org/10.1145/2514601.2514624>
- Conrad JG, Zeleznikow J (2015) The role of evaluation in AI and law: An examination of its different forms in the AI and law journal. *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, 181–186
- Craandijk D, Bex F (2021) Deep learning for abstract argumentation semantics. *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 1667–1673
- Fest I, Meijer A, Schäfer M, van Dijk J (2024) Values? Camera? Action! An ethnography of an AI camera system used by the Netherlands Police. *Under Review*
- Francesconi E (2022) The winter, the summer and the summer dream of artificial intelligence in law. *Artif Intell Law* 30(2):147–161. <https://doi.org/10.1007/s10506-022-09309-8>
- Fratrič P, Parizi MM, Sileno G, van Engers T, Klous S (2023) Do agents dream of abiding by the rules? Learning norms via behavioral exploration and sparse human supervision. *Proc Nineteenth Int Conf Artif Intell Law* 81–90. <https://doi.org/10.1145/3594536.3595153>

- Fuchs M, Jadhav A, Jaishankar A, Cauffman C, Spanakis G (2023) What's wrong with this product?': detection of product safety issues based on information consumers share online. *Proc Nineteenth Int Conf Artif Intell Law* 397:401. <https://doi.org/10.1145/3594536.3595145>
- Gan L, Kuang K, Yang Y, Wu F (2021) Judgment prediction via injecting legal knowledge into neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14), 12866–12874
- Gray M, Savelka J, Oliver W, Ashley K (2023) Automatic Identification and Empirical Analysis of Legally Relevant Factors. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 101–110. <https://doi.org/10.1145/3594536.3595157>
- Habba E, Keydar R, Bareket D, Stanovsky G (2023) The Perfect Victim: Computational Analysis of Judicial Attitudes towards Victims of Sexual Violence. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 111–120. <https://doi.org/10.1145/3594536.3595168>
- Haim A, Kesari A (2023) Image Analysis Approach to Trademark congestion and depletion. *Proc Nineteenth Int Conf Artif Intell Law* 402–406. <https://doi.org/10.1145/3594536.3595126>
- Herrewijnen E, Nguyen D, Mense J, Bex F (2021) Machine-annotated Rationales: Faithfully Explaining Text Classification. *Workshop on Explainable Agency in Artificial Intelligence*. AAAI 2021
- High R (2012) The era of cognitive systems: an inside look at IBM Watson and how it works. IBM Corporation Redbooks 1:16
- Hillebrand L, Pielka M, Leonhard D, Deußer T, Dilmaghani T, Kliem B, Loitz R, Morad M, Temath C, Bell T, Stenzel R, Sifa R (2023) sustain.AI: A Recommender System to analyze Sustainability Reports. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 412–416. <https://doi.org/10.1145/3594536.3595131>
- Hultstijn J (2023) Computational accountability. *Proc Nineteenth Int Conf Artif Intell Law* 121–130. <https://doi.org/10.1145/3594536.3595122>
- Jiang C, Yang X (2023) Legal syllogism prompting: teaching large Language models for Legal Judgment Prediction. *Proc Nineteenth Int Conf Artif Intell Law* 417–421. <https://doi.org/10.1145/3594536.3595170>
- Katz DM, Bommarito MJ, Gao S, Arredondo P (2023) Gpt-4 passes the bar exam. Available SSRN 4389233
- Lawrence C, Cui I, Ho D (2023) The bureaucratic challenge to AI governance: an empirical Assessment of implementation at U.S. Federal agencies. *Proc 2023 AAAI/ACM Conf AI Ethics Soc* 606–652. <https://doi.org/10.1145/3600211.3604701>
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521(7553):436–444
- Li O, Liu H, Chen C, Rudin C (2018) Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1)
- Lundberg SM, Lee S-I (2017) A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst*, 30
- Marcus G, Davis E (2019) *Rebooting AI: building artificial intelligence we can trust*. Vintage
- Mumford J, Atkinson K, Bench-Capon T (2023) Combining a legal knowledge model with machine learning for reasoning with legal cases. *Proc Nineteenth Int Conf Artif Intell Law* 167–176. <https://doi.org/10.1145/3594536.3595158>
- Nielsen A, Skylaki S, Norkute M, Stremitzer A (2023) Effects of XAI on legal process. *Proc Nineteenth Int Conf Artif Intell Law* 442–446. <https://doi.org/10.1145/3594536.3595128>
- Nieuwenhuizen E, Meijer A, Bex F, Grimmelikhuisen S (2023) Explanations increase citizen trust in police algorithmic recommender systems: findings from two experimental tests. Under Rev
- Odekerken D, Bex F, Borg A, Testerink B (2022) Approximating stability for applied argument-based inquiry. *Intell Syst Appl* 16:200110. <https://doi.org/10.1016/j.iswa.2022.200110>
- Odekerken D, Bex F, Prakken H (2023) Justification, stability and relevance for case-based reasoning with incomplete focus cases. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 177–186
- Pasquale F (2015) *The black box society: the secret algorithms that control money and information*. Harvard University Press
- Peters JG, Bex FJ, Prakken H (2023) Model-and data-agnostic justifications with A Fortiori Case-Based Argumentation. *19th International Conference on Artificial Intelligence and Law*, 207–216
- Piccolo SA, Katsikouli P, Gammeltoft-Hansen T, Slaats T (2023) On predicting and explaining asylum adjudication. *Proc Nineteenth Int Conf Artif Intell Law* 217–226. <https://doi.org/10.1145/3594536.3595155>
- Power M (1994) *The audit explosion* (Issue 7). Demos
- Prakken H (2010) An abstract framework for argumentation with structured arguments. *Argument Comput* 1(2):93–124

- Prakken H, Sartor G (2015) Law and logic: a review from an argumentation perspective. *Artif Intell* 227:214–245. <https://doi.org/10.1016/j.artint.2015.06.005>
- Ribeiro MT, Singh S, Guestrin C (2016) ‘Why should i trust you?’ Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144
- Riera J, Solans D, Karimi-Haghighi M, Castillo C, Calsamiglia C (2023) Gender Disparities in Child Custody Sentencing in Spain: A Data Driven Analysis. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 237–246. <https://doi.org/10.1145/3594536.3595135>
- Rissland EL, Ashley KD, Branting LK (2005) Case-based reasoning and law. *Knowl Eng Rev* 20(3):293–298
- Robeer M, Bex F, Feelders A (2021) Generating realistic natural language counterfactuals. *Find Association Comput Linguistics: EMNLP 2021*, 3611–3625
- Robeer M, Bron M, Herrewijnen E, Hoeseni R, Bex F (2023) The Explabox: Model-Agnostic ML Transparency & Analysis. *Under Review*
- Sætra HS (2022) AI for the sustainable development goals. *CRC*
- Santin P, Grundler G, Galassi A, Galli F, Lagioia F, Palmieri E, Ruggeri F, Sartor G, Torroni P (2023) Argumentation structure prediction in CJEU decisions on fiscal state aid. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 247–256
- Sarker MK, Zhou L, Eberhart A, Hitzler P (2021) Neuro-symbolic artificial intelligence. *AI Commun* 34(3):197–209
- Savelka J (2023) Unlocking Practical Applications in Legal Domain: Evaluation of GPT for Zero-Shot Semantic Annotation of Legal Texts. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 447–451. <https://doi.org/10.1145/3594536.3595161>
- Schirmer M, Nolasco IMO, Mosca E, Xu S, Pfeffer J (2023) Uncovering Trauma in Genocide Tribunals: An NLP Approach Using the Genocide Transcript Corpus. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 257–266. <https://doi.org/10.1145/3594536.3595147>
- Schraagen M, Bex F (2019) Extraction of semantic relations in noisy user-generated law enforcement data. *Proceedings of the 13th IEEE International Conference on Semantic Computing (ICSC 2019)*
- Schraagen M, Brinkhuis M, Bex F (2017) Evaluation of Named Entity Recognition in Dutch online criminal complaints. *Comput Linguistics Neth J* 7:3–16
- Selten F, Robeer M, Grimmelikhuijsen S (2023) Just like I thought*: Street-level bureaucrats trust AI recommendations if they confirm their professional judgment. *Public Adm Rev* 83(2):263–278
- Servantez S, Lipka N, Siu A, Aggarwal M, Krishnamurthy B, Garimella A, Hammond K, Jain R (2023) Computable Contracts by Extracting Obligation Logic Graphs. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 267–276
- Steenhuis Q, Willey B, Colarusso D (2023) Beyond Readability with RateMyPDF: A Combined Rule-based and Machine Learning Approach to Improving Court Forms. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 287–296
- Stern RE, Liebman BL, Roberts ME, Wang AZ (2020) Automating fairness? Artificial intelligence in the Chinese courts. *Colum J Transnat'l L* 59:515
- Tan H, Zhang B, Zhang H, Li R (2020) The sentencing-element-aware model for explainable term-of-penalty prediction. *Natural Language Processing and Chinese Computing: 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14–18, 2020, Proceedings, Part II* 9, 16–27
- Tolan S, Miron M, Gómez E, Castillo C (2019) Why machine learning may lead to unfairness: Evidence from risk assessment for juvenile justice in catalonia. *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, 83–92
- Unver MB (2023) Rebuilding ‘ethics’ to govern AI: How to re-set the boundaries for the legal sector? *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 306–315. <https://doi.org/10.1145/3594536.3595156>
- van Droffelaar IS, Kwakkel JH, Mense JP, Verbraeck A (2022) Simulation-Optimization Configurations for Real-Time Decision-Making in Fugitive Interception. *Available at SSRN 4659539*
- van Leeuwen L, Verheij B, Verbrugge R, Renooij S (2023) Using Agent-Based Simulations to Evaluate Bayesian Networks for Criminal Scenarios. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 323–332. <https://doi.org/10.1145/3594536.3595125>
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is All you Need. *Advances in Neural Information Processing Systems*, 30. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fdb053c1c4a845aa-Abstract.html>

-
- Verheij B (2020) Artificial intelligence as law. *Artif Intell Law* 28(2). <https://doi.org/10.1007/s10506-020-09266-0>
- Westermann H, Benyekhlef K (2023) JusticeBot: A Methodology for Building Augmented Intelligence Tools for Laypeople to Increase Access to Justice. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 351–360. <https://doi.org/10.1145/3594536.3595166>
- Zhang G, Nulty P, Lillis D (2023) Argument Mining with Graph Representation Learning. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 371–380
- Ziewitz M (2016) Governing algorithms: myth, mess, and methods. *Sci Technol Hum Values* 41(1):3–16
- Zin MM, Nguyen HT, Satoh K, Sugawara S, Nishino F (2023) Improving translation of case descriptions into logical fact formulas using legalcasener. *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 462–466

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.