JOURNAL OF
CROSS-DISCIPLINARY RESEARCH
IN COMPUTATIONAL LAW [CRCL]

Transdisciplinary research as a way forward in AI & Law

Floris Bex*

Abstract

The field of Artificial Intelligence & Law is a community of law and computer science scholars, with a focus on AI applications for the law and law enforcement. Such applications have become the subject of much debate, with techno-pessimists on one side and techno-optimists on the other. What is the role of the AI & Law community in this debate, and how can we investigate AI for the law without losing ourselves in the "algorithmic drama" of overly optimistic versus strongly pessimistic camps? I will argue for three points: (1) combine research on data-driven systems, such as generative AI, with research on knowledge-based AI; (2) put AI into (legal) practice, working together with courts, the police, law firms and citizens; (3) work together across disciplines, bringing together those who think about how to build AI and those who think about how to govern and regulate it.

Keywords: AI & Law; combining data-driven and knowledge-based systems; AI in legal practice; interdisciplinary research

Journal of Cross-disciplinary Research in Computational Law

DOI: pending

www.journalcrcl.org

* Associate Professor at Utrecht University & Professor at Tilburg University, f.j.bex@uu.nl

Introduction

We live in exciting times for AI & Law: deep learning, and particularly generative language models, can finally deliver on AI's promises, with GPT-4 passing the US bar exam [Katz et al. 2024]; many academic papers on legal language processing and its applications are being published [Gan et al. 2021; Jiang and Yang 2023; Zhong et al. 2019; Zhou et al. 2023]; governments are investing in legal NLP to make the law more efficient, consistent and accessible;[1] and the Legal Tech sector is on the rise, with companies developing a wide array of (NLP/ GPT) products for legal services.[2] Another perspective however, paints a more negative picture. The openness[3] and trustworthiness[4] of models like GPT are being questioned. The hype surrounding 'robot judges' has led to debate on the nature and ethics of such automated legal decision-making [Araujo et al. 2020; Bex 2024; Pasquale and Cashwell 2018]. More generally, many prominent researchers and opinion makers are warning about the dangers of generative AI systems.[5]

All this positive and negative excitement has drawn us further into Ziewitz' "algorithmic drama" [Ziewitz 2016], where AI algorithms determine a large part of our lives and societies, algorithms that are complex, opaque and outside of our control. The AI & Law research community should analyse and investigate AI for the law and how AI influences the law without getting caught up in this drama, steering clear of both excessive techno-optimism or techno-pessimism. But how can we do that? I will argue that three points are important for a way forward in AI & Law:

(1) combining knowledge & data in AI;

(2) evaluating how AI & Law is used in practice; and

(3) combining different disciplines: law, AI and beyond.

In this position paper,[6] I will discuss the three points in the context of the field of AI & Law, showing where it is ahead, and where it is behind. I will further give two example cases of research I have been involved in where we have tackled these points head-on.

Artificial Intelligence and Law's role in the algorithmic drama

The field of AI & Law is a largely "techno-optimistic" community of law and computer science scholars, with a focus on AI applications for the law and legal reasoning, and the AI techniques underlying these applications. AI & Law's main technical focus – *AI for Law* - sets it apart from the Law & Technology community, which has a legal focus – *Law for AI*. That said, historically there has always been a natural overlap between both communities,[7] and the AI & Law community also increasingly addresses, for example, the legal implications of the use of AI applications in the law.

A large part of the international community is centered around the International Association for AI and Law,[8] the biennial International Conference on Artificial Intelligence

[1] 'China wants legal sector to be AI-powered by 2025', ZDnet, 12 December 2022 (`https://www.zdnet.com/article/china-wants-legal-sector-to-be-ai-powered-by-2025/`).

[2] The total number of companies in the CodeX Techindex has risen from just over 700 in 2017 to almost 2200 in 2023 (`https://techindex.law.stanford.edu/`).

[3] 'OpenAI's GPT-4 Is Closed Source and Shrouded in Secrecy', Vice.com, 16 March 2023 (`https://www.vice.com/en/article/ak3w5a/openais-gpt-4-is-closed-source-and-shrouded-in-secrecy`)

[4] 'ChatGPT: US lawyer admits using AI for case research', BBC News, 27 May 2023 (`https://www.bbc.com/news/world-us-canada-65735769`).

[5] 'Pause Giant AI Experiments: An Open Letter', Future of Life Institute, 22 March 2023 (`https://futureoflife.org/open-letter/pause-giant-ai-experiments/`).

[6] This paper can be considered a companion paper to [Bex 2024], which is based on my Presidential Address at ICAIL 2023. The current paper goes into more depth regarding the three different points and discusses different case studies.

[7] At the first ICAIL conference in 1987, the number of first authors from law schools and computer science departments was roughly 50-50, and the early Bileta conferences (one of the oldest Law & Technology conferences) included technical articles on, e.g., legal expert systems. (`https://www.bileta.org.uk/`)

[8] `http://iaail.org/`.

[9] See `http://iaail.org/?q=page/past-icails`. For a historical overviews of the ICAIL conferences up to 2011, see [Bench-Capon et al. 2012]

and Law (ICAIL).[9] and the journal Artificial Intelligence and Law.[10] There is also the annual Jurix conference on Legal Knowledge and Information Systems,[11] the Jurisin series of workshops on Juris-informatics,[12] and a number of related communities, workshops and journals.[13]

In this section, I will give a brief and non-exhaustive overview of how the core AI & Law community has been engaging with the above-mentioned three points.

Combining Knowledge-based and Data-driven AI

In Artificial Intelligence there are two dominant approaches. The first is symbolic, knowledge-based AI, where algorithms reason based on pre-programmed knowledge codified into, for example, rules. The second is machine learning, data-driven AI that learns to recognize (complex) patterns given large amounts of data.[14]

AI & Law has for most of its history been dominated by the first approach, looking mainly at knowledge-based, logical models of legal reasoning like argumentation [Prakken and Sartor 2015] and case-based reasoning [Rissland et al. 2005]. Early on, legal expert systems based on these logical models could count on much interest,[15] but in the 1990s the interest diminished as they could not live up to the hype [Leith 2010]. It is difficult and time-consuming to include all the relevant (legal) expert knowledge in a system,[16] and symbolic, knowledge-based systems are notoriously bad at handling noisy or ambiguous input such as natural language or open-textured legal concepts. However, knowledge-based approaches still make sense in bounded domains, where the law is relatively simple, static and known and where rule-based systems have become commonplace, with everything from basic HR policies to tax law being encoded in "business rules" that are used to infer, calculate and decide on legal conclusions. Furthermore, systems with explicit rules are more readily interpretable than the complex neural networks that are used in modern machine learning.

Since 2015, data-driven (deep) machine learning has become the dominant approach in AI & Law. Advances in Natural Language Processing (NLP) spurred new talk of 'robo-judges' that can predict outcomes of legal cases,[17] and countries such as China are encouraging and deploying such systems in court.[18] NLP has also been used for applications such as finding similar cases [Dan et al. 2023; Mandal et al. 2021], legal search [Custis et al. 2019; Dadgostari et al. 2021] and legal text summarization [Deroy et al. 2024; Schraagen et al. 2022; Zhong et al. 2019]. The core problem with these data-driven approaches is that they do not perform or understand legal reasoning, but instead work with correlations between pieces of text (words, sentences) in the training data. Furthermore, machine learning models are often hard to interpret, and they cannot provide explanations for their behaviour in terms of, for example, legal rules in the way that knowledge-based systems can.[19]

[10] See `https://www.springer.com/journal/10506`. For a historical overview of the AI & Law journal, see the special issue on *Thirty Years of Artificial Intelligence and Law* [Bench-Capon 2022].

[11] See `http://jurix.nl/`.

[12] See `https://research.nii.ac.jp/~ksatoh/jurisin2023/`.

[13] For example, the Codex FutureLaw conference (`https://conferences.law.stanford.edu/futurelaw/`), workshops like NLLP (`https://nllpw.org/`), and of course in the CRCL journal and conference series.

[14] Surden [Surden 2019] speaks of *Machine Learning* and *Rules, Logic and Knowledge Representation*, Hildebrandt and Diver [Diver 2021; Hildebrandt 2018] distinguish between *data-driven law* and *code-driven law*, cf. e.g. (`https://www.cohubicol.com/`).

[15] For example, McCarthy's TAXMAN system [McCarty 1977]. See also more generally Susskind [Susskind 1986] on legal expert systems.

[16] The so-called '*knowledge acquisition bottleneck*', see e.g., [Cullen and Bryman 1988].

[17] The discussion started with the seminal work by Aletras et al. [2016]. See, e.g., [Babic et al. 2020] for a positive assessment of legal prediction, [Pasquale and Cashwell 2018] for a critique on legal prediction, and [Bex and Prakken 2021] for a discussion on the relevance of such predictive algorithms.

[18] See [Stern et al. 2020] for descriptions of AI projects in Chinese courts and references to the original Chinese sources.

[19] There is work on language models like GPT performing legal reasoning and explaining themselves, see, e.g., [Guha et al. 2023; Katz et al. 2024; Savelka 2023]. However, whether these models can reason is debatable – for example, they are not (yet) very good at applying legal rules or statutes [Blair-Stanek et al. 2023; Guha et al. 2023]. Furthermore, language models have problems with their consistency and correctness, being prone to so-called "hallucinations", where (linguistically) plausible but (legally, factually) incorrect text is being generated.

In the search for transparent and scalable AI that can perform both legal reasoning and handle noisy and open-textured concepts, it has been argued that data and knowledge – or learning and reasoning – should be combined in a single "neuro-symbolic" system [Marcus and Davis 2019; Sarker et al. 2021].[20] The most common example of such a type of hybrid learning/reasoning system in AI & Law is when data-driven machine learning techniques are used to extract legal knowledge – rules, arguments, cases – from unstructured data such as text, and knowledge-based reasoning techniques are subsequently used to reason with this information (e.g., [Ashley and Walker 2013; Brüninghaus and Ashley 2005; Mumford et al. 2023; Odekerken et al. 2022; Schraagen et al. 2018]. Another type of hybrid system that has been proposed in AI & Law is one where knowledge-based techniques are used to explain the outcomes of data-driven machine learning (e.g., [Čyras et al. 2019; Grabmair 2017; Peters et al. 2023; Prakken and Ratsma 2022]). In these hybrid systems, the reasoning is still done by the knowledge-based part. In other approaches to integrating learning and reasoning, which are still quite uncommon in AI & Law, the system is a pure machine learning system, but it is used to solve typical knowledge-based problems, like argumentation or case-based reasoning (e.g., [Craandijk and Bex 2021, 2022; Li et al. 2018; Otero et al. 2023]). The advantage of such "neural solvers" is that they can output not just a conclusion, but also the reasoning that led to that decision. Other authors have also trained pure machine learning systems, but they have used knowledge to constrain what the system can learn so that, for example, it will follow legal rules or base its predictions on legally relevant pieces of text (e.g., [Gan et al. 2021; Santosh et al. 2022; Zhou et al. 2023]).

So, while the data-driven approach has very much become the dominant approach in AI & Law, we see that the knowledge-based approach still is – and should be - influential. Especially in the law, it is important that decisions are based on more than just correlations, and that decisions are made in ways that are transparent, contestable and in line with the law.

Evaluating AI & Law in practice

In a multi-faceted field such as AI & Law, there are many types of research with at least as many types of evaluation. For research into, for example, logical models of legal reasoning we have evaluation by means of mathematical theorems and their formal proofs, but also evaluation in terms of how well such models 'fit' with the (legal) reality.[21] While for the latter one can make a philosophical "argument from intuition",[22] it is also possible to more empirically test different types of models by performing larger case studies.[23] Then there is empirical evaluation of data-driven models by, for example, making (statistical) comparisons between different models, and between model outputs and some gold standard dataset. Lately there has also been work that by itself is more evaluative in nature, where a specific AI technology is subjected to a legal evaluation [Guo and Kennedy 2023], or the impact of AI technology on legal reasoning is evaluated empirically [Barysė and Sarel 2024].

In addition to research into computational models of the law, AI & Law also has a focus on innovative AI applications for the legal domain. Much of the research on the more data-driven systems is already application-oriented in the sense that it presents a solution for a certain legal task, such as legal document summarization or legal search. As mentioned above, the evaluation of such systems is mainly performed by means of various (often quantitative/statistical) comparisons to other systems or gold standard datasets, with human users or experts only being considered in some studies (e.g., [Habernal et al. 2023]). Researchers in AI & Law have also been working on (prototype) applications for legal practice. There are various ways to evaluate such systems. For any kind of real application, it is important to perform usability tests, but

[20] Note that the data-driven components of such a system need to be based on neural networks per se – they can also use more traditional (non-neural) machine learning algorithms.

[21] For example, case-based reasoning models the type of reasoning often seen in common law, and argumentation based on rules models the reasoning with laws and statutes that is more common in continental law.

[22] See '*Intuition*' in the Stanford Encyclopedia of Philosophy (`https://plato.stanford.edu/entries/intuition/`).

[23] See, e.g., [Prakken et al. 2020], in which multiple (formal) models for reasoning with proof are compared by modelling the same murder case in all the different models.

also to evaluate the (kind of) impact it has on legal decision making. Furthermore, a broader evaluation on the legal, organisational and societal implications of using an AI application in the legal field is also desirable. In 2015, Conrad and Zeleznikow performed a structured analysis on the types of evaluation included in articles about an application in the AI & Law journal [Conrad and Zeleznikow 2015]. They found that an operational-usability evaluation was only presented in 7% of all articles. An informal analysis of journal articles from 2023 shows that this has improved significantly: at least 30% of the articles presenting an application have a form of (expert) user evaluation or involvement, with some even evaluating their application in practical setting, with actual users [Bakhshayesh and Abbasianjahromi 2023; Lettieri et al. 2023; Marković and Gostojić 2023].

So, the evaluation of AI & Law systems with users has been steadily increasing. However, the number of AI & Law applications that have resulted in practical applications is still relatively small.[24] While this is partly to be expected — it is not our goal as researchers to build tools for industry or the legal field — working and evaluating with stakeholders from practice is necessary if we really want to find out the impact of AI & Law on legal decision making, organisations and society.

Working with different disciplines

AI & Law is a technically focused community mainly made up of computer science and tech-savvy law scholars, with connections to the broader legal field. Over the years, the interest of the legal community in AI & Law has waxed and waned together with the general societal interest for and the use of AI. After the "fall of the expert system" [Leith 2010] in the 1990s, many legal practitioners and scholars lost interest in AI, switching over to the legal study of other kinds of technology such as the internet. With the renewed interest in data (protection) and AI after 2015, we also see the interest of the legal community for AI & Law return.[25]

This has led to an increasing number of more legally focused articles at AI & law conferences and in journals: on "law-by-design" [Almada 2019] – how legal concepts can be directly implemented in AI systems – on legal aspects of AI for the legal sector [Guo and Kennedy 2023; Unver 2023], and on the effects of AI on the legal process [Nielsen et al. 2023]. Furthermore, there has also been an increase in research on empirical legal studies using NLP [Chandler et al. 2023; Habba et al. 2023; Piccolo et al. 2023; Riera et al. 2023; Schirmer et al. 2023], and the use of AI techniques such as agent-based simulations [Leeuwen et al. 2023] and machine learning [Fratrič et al. 2023] to study how the law and law enforcement work.

With AI becoming commonplace in today's society, however, the application of AI to the law is no longer just of interest to lawyers and computer scientists. This is also evident when looking at the traditional AI & Law venues: in the past few years alone, we have seen publications by researchers from diverse fields such as social studies [Witt et al. 2023], management [Cohen et al. 2023; Yalcin et al. 2023], business [Braun 2023; Yalcin et al. 2023] and administrative sciences [Saragih et al. 2023], criminology [Simmler et al. 2023], economics [Di Porto 2023], information management [Lawrence et al. 2023; McLachlan et al. 2023] and psychology [Barysė and Sarel 2024]. Looking at the broader picture, we also see researchers from, for example, philosophy [Allo 2022], ethics [Pruss 2023], public administration [Nieuwenhuizen et al. 2023; Soares et al. 2023], and communication studies [Araujo et al. 2020] discussing AI and the law. Each of these disciplines brings its own insights: people from management or information sciences allow us to zoom out and see the bigger socio-technical systems surrounding the technology, and researchers from psychology and public administration look at our technology through an empirical lens. More critical humanities, such as philosophy, ethics, but also communication and media studies question some of the core behaviours and ways of communicating that the AI & Law community takes for granted.

[24] Examples of such applications can be found in [Odekerken et al. 2022], which has been implemented at the Dutch police (see Section "*AI for citizen complaint intake at the Dutch police*"), and [Al-Abdulkarim et al. 2019], which has been used at law firms in the UK. Further examples are [Cohen et al. 2023; Westermann and Benyekhlef 2023], who have started openly accessible platforms or websites for their services.

[25] At ICAIL 2021, the number of first authors from law schools and computer science departments was roughly 30-70. There were quite a few articles which were co-authored by researchers from law and computer science, as well as authors that work at both law schools and computer science departments.

So, while the core AI & Law community is still largely composed of more technical, computer science minded researchers, Artificial Intelligence and Law in the broader sense is steadily gaining interest from more academic disciplines.

Putting the three points into action: examples of two projects

We see that the AI & law community has changed quite a bit since its inception in the 1980s: there is more focus on data-driven and hybrid systems, and more innovative applications are being developed and evaluated from different disciplinary perspectives.

I will now provide two examples of projects that I believe are exemplary for the changing field of AI & Law, in which we aim to tackle the three points head on.

AI for citizen complaint intake at the Dutch police

The first project was done in the context of the National Police Lab AI,[26] and concerns an AI system for the intake of citizen complaints about online trade fraud; for example, false web shops or malicious traders on eBay not delivering products to people. The police receive more than 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. To save the police from having to manually check all the reports, we developed a recommender system that determines whether a case is possibly fraud based on the submitted complaint form, and it only recommends filing an official report if the recommender system considers it fraud. This system was implemented by the police in 2018, and it is still being used in 2023.

For the intake system, we combined knowledge- and data-driven AI. The system has a legal model of the domain that captures the relevant part of the Dutch Criminal Code and police policy rules in a rule-based argumentation model. Examples of rules in this model are "**if** the product was paid for but not sent **and** deception was used, **then** it is possibly fraud" and "**if** the supposed seller used a false location **or** a false website **then** deception was used" [Borg and Bex 2021; Odekerken et al. 2022]. Because the complaint reporting form also contains a free text field where citizens can tell their story, the system also includes natural language processing techniques to extract basic citizen observations like "I paid but did not receive anything" from the text field.[27] Using the basic observations and the legal rule model, the system then tries to infer whether the complaint is possibly a case of fraud or not. If it turns out there are still missing observations that can change the system's conclusion, the system can ask the citizen for these observations. Determining whether the conclusion could still change, and which observations are still relevant for such a change, is computationally quite expensive, meaning that it takes a standard algorithm several minutes to do this. Because we do not want the citizen to have to wait this long, we developed and algorithm that can do it almost instantaneously [Odekerken et al. 2022]. Once all relevant questions have been asked, the system will present the recommendation whether to file an official complaint or not to the user, together with an explanation for its recommendation in terms of the (legal) rules and observations it used to infer the conclusion. If the user decides to file an official complaint, the input thus far is transferred to a secure environment where the citizen can electronically sign the complaint and send it to the police.

The intake system has been evaluated internally by the police on various aspects, such as accuracy, user satisfaction, efficiency, and effectiveness. The system turned out to be between 80% and 90% accurate when measured against what police case workers would recommend (submit or do not submit report), and the efficiency of the reporting process was increased significantly. We further performed

[26] National Police lab AI is a collaboration between multiple universities and the Netherlands National Police, where many of the PhDs also work for or at the Police, combining research and development of AI (`https://www.uu.nl/onderzoek/ai-labs/nationaal-politielab-ai`).

[27] We experimented with various machine learning NLP approaches to extract entities and (event) relations from the text [Schraagen and Bex 2019; Schraagen, Brinkhuis, et al. 2017]. While results were acceptable, the final implementation depends on regular expressions to extract observations.

two more scientific evaluations of the system. The first of these evaluations [Nieuwenhuizen et al. 2023] concerned the system's effect on human trust: would citizens mind receiving 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. In the experiment, the system told the participants it was probably *not* criminal fraud in their case, and therefore recommended not to file an official report. We then measured the participants' trusting behaviour: did they still file an official report? 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. Of the group that did receive an explanation, 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.[28]

The second type of evaluation concerned an ethnographic case study at the department of the Dutch Police that processes the incoming citizen reports on (alleged) online trade fraud [Soares et al. 2023]. Before the intake system was implemented, human case workers had to manually get the observations from the free text of the intake form, and email with the complainant to ask any further questions to complete the file. These manual tasks were taken over by the system, and the system thus provides the case workers with more complete and structured information. Furthermore, the system also provided the case workers with the recommendation it had given the citizen (i.e., possibly fraud – submit a report; or probably not fraud – do not submit a report), although it did not provide an explanation for this.[29] Note that it was still up to the case worker to decide on a case-by-case basis whether the original report plus the citizen's answers would be entered into the police systems as a fraud report. What we observed was that the system helped the case workers by providing structured data, hence allow them to focus more on assessing cases that needed a nuanced judgement. However, the recommendation of the system (fraud or no fraud) was simply ignored by the case workers – one of the key reasons for this was that the system did not provide an explanation or rationale for its recommendations.

This shows how AI can be designed and evaluated in practice from different perspectives, using different methods: computational, experimental and ethnographic. We see that explanations are important for citizen trust,[30] and only a knowledge-based system that includes explicit (legal) rules can provide such explanations. On the other hand, more data-driven natural language processing is also needed to allow citizens to interact with the system in a natural way. Furthermore, screen-level bureaucrats, like police case workers, were happy to use the system as an assistant because their professional discretion was not threatened by it.

AI for supporting paralegals with traffic fine appeals cases at a Dutch court.

In the second project, we performed a critical case study of the development and use of an AI decision support system for processing traffic violation appeals at a Dutch court [Kolkman et al. 2023]. For minor traffic fines such as speeding or parking violations, the public prosecutor in the Netherlands hands out an administrative fine. If the person concerned does not agree with this fine, an appeal can be made first with the public prosecution and then with the court – the latter we examined in our case study. There are more than 50 such appeal cases per week at a single court, and paralegals spend on average 20 minutes preparing a case, so aiding the paralegals in this administrative decision-making process could alleviate the workload significantly.

[28] Internal evaluations at the police showed that (with an explanation) about 60% still filed the report even though they were recommended not to. There was some evidence that for actual complaints the lower trusting behaviour was caused by the fact that people were angrier and more frustrated because they had lost actual money, and they wanted the police's help in getting it back even if they were told their case was clearly not one of criminal fraud.

[29] Note that because the intake system was originally not designed for the case workers, this explanation could only be given in the citizen interface and not in the interface the case workers worked with.

[30] And possibly also for police case worker trust since they ignored the recommendations without an explanation.

As there were no AI applications in the courts when we started our research, we worked together with the court[31] to develop a system to help paralegals prepare appeal cases for court hearings. The paralegals can upload a PDF file of the appeal case to the system, which first uses basic language processing to extract a structured case overview – case number, type of offence, brief description, height of the fine, and various deadlines such as those for payment of the fine and lodging of the appeal. It then automatically checks whether the deadlines have been met, showing this to the paralegal. The system then, based on the appellant's and prosecution's arguments, identifies similar cases using more advanced NLP techniques such as document vectorization and comparison. It presents these similar case documents as a list of search results that can be clicked on to view similar cases. Finally, the system suggests the most likely outcome of the case – affirmed, rejected, inadmissible, or modify the decision of the prosecution (e.g., a lower fine) – using legal text classification.[32] Because the system does not arrive at the recommendation by following legal rules like in the fraud intake system, the system cannot provide an explanation or rationale why it recommends an outcome.

We developed and tested the system extensively with three paralegals of a Dutch court and made some interesting findings. First, the paralegals indicated that the automatic extraction, overview and checking of information from the free text of the case document was very useful, as it saved them from having to look this up in the original document. Second, the similar-case matching was also useful, as it allowed paralegals to search all the previous cases that were in the database. When questioned about whether the similar-case-matching algorithm could lead to biased results, the paralegals indicated that in their current process each paralegal only has access to previous cases they themselves have handled, so they are already biased towards their own previous cases. Finally, the suggestion of the most likely outcome was deemed useless without an explanation – the paralegals went and read the various arguments in the appeal themselves for a professional verdict.

This project, like the one with the police, demonstrates the value of actively working with practice in building and evaluating systems. For example, it becomes possible to study the actual use of algorithmic systems in organisations that are not normally known to use such systems, such as courts. Interestingly, our findings at the court largely coincide with what we found for the other screen-level bureaucrats, the police case workers: relatively basic systems that structure and gather information are seen as a positive thing, as long as these systems do not impinge on the bureaucrats' professional discretion as decision makers. Furthermore, recommendations or predictions which are not backed up by an explanation or rationale are ignored.

Conclusion

Artificial Intelligence and Law is an interdisciplinary, techno-optimistic community with a long history going back at least as far as the legal expert systems in the 1980s. Like any community interested in computational law, advancements in modern AI require the AI & Law community to think about its future and where it wants to position itself in today's 'algorithmic drama'. I have argued that three points are important when thinking about a way forward in AI & Law: (1) combining knowledge & data in AI; (2) evaluating how AI & Law is used in practice; and (3) combining different disciplines.

When looking at where AI & Law is now regarding point (1), we see that there is a steadily increasing amount of work on hybrid systems that combine separate data-driven learning modules and knowledge-based reasoning modules, but less work on true "neuro-symbolic" systems where modern machine learning techniques are used for reasoning. With respect to point (2), more AI applications are being researched and developed for the legal field, and evaluations with (real) users are also being increasingly performed, but actual practical applications based on core AI & Law research that are being used and evaluated in the legal field

[31] Thus, our case study had a strong action research and participatory component, cf. [Davison et al. 2004].

[32] That is, a machine learned legal case prediction algorithm similar to [Aletras et al. 2016].

are still quite scarce.[33] Finally (point 3), even though the core field is still made up out of more technically minded researchers, other scholars from a wide variety of academic disciplines are becoming interested in AI & Law.

Looking at two recent projects, we can see that they are exemplary for the path the field of AI & Law, in my view, should be taking on the three points. The online trade fraud complaint intake system from Section "*AI for citizen complaint intake at the Dutch police*" demonstrates the combination of knowledge and data, and the advantages that brings with respect to transparency. Whereas both systems – the intake system at the police and the support system at the court – use data-driven natural language processing to extract information from text, only the intake system models the actual legal domain rules. So only the intake system can provide meaningful explanations for its recommendations in terms of such legal rules. With respect to evaluation in practice, we see in both cases that screen-level bureaucrats like police case workers or paralegals want to remain in control, and do not follow recommendations blindly without an explanation. This runs contrary to both the perceived usefulness and dangers of AI systems that predict or recommend an outcome in a case. Finally, the two cases show how AI can be designed and evaluated in practice from different disciplinary perspectives, using different methods: computational, experimental, ethnographic, participatory. We have worked with researchers from computer and data science, law, public management and media studies, without whom we would not have been able to design, build and evaluate the systems in the manner that we did.

Today's AI & Law community includes many disciplines and stakeholders, and studies different types of AI for Law and Law for AI in a broad societal framework. Even though data-driven machine learning has led to impressive advances in AI recently, we should not consider it to be the only kind of AI - having the next generation of large language models take the law into account will require insights and techniques from knowledge-based approaches. These hybrid data/knowledge systems should be evaluated with stakeholders from practice - we cannot rightly claim to develop AI for the legal field, if ultimately only very few in that field can or will use our systems and techniques, or at least derivatives of them. And finally, a mature AI & Law field should look beyond just (legally informed) Computer Science to other disciplines, considering the bigger socio-technical systems surrounding the technology, and questioning the core concepts the AI & Law community takes for granted. Only with a diverse palette of researchers and methods can we responsibly design and analyse the future AI for the law.

References

Latifa Al-Abdulkarim, Katie Atkinson, Trevor Bench-Capon, Stuart Whittle, Rob Williams, and Catriona Wolfenden. 2019. 'Noise induced hearing loss: Building an application using the ANGELIC methodology'. *Argument & Computation*, 10, 1, 5–22.

Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoţiuc-Pietro, and Vasileios Lampos. 2016. 'Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective'. *PeerJ computer science*, 2, e93.

Patrick Allo. 2022. 'The Interpretability Problem and the Critique of Abstraction (s)'. *Journal of Cross-Disciplinary Research in Computational Law*, 1, 2.

Marco Almada. 2019. 'Human intervention in automated decision-making: Toward the construction of contestable systems'. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law* (ICAIL '19). Association for Computing Machinery, New York, NY, USA, 2–11.

Theo Araujo, Natali Helberger, Sanne Kruikemeier, and Claes H. de Vreese. 2020. 'In AI we trust? Perceptions about automated decision-making by artificial intelligence'. *AI & SOCIETY*, 35, 3, 611–623.

Kevin D. Ashley and Vern R. Walker. 2013. 'Toward constructing evidence-based legal arguments using legal decision documents and machine learning'. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*. ACM, Rome Italy, 176–180.

Boris Babic, Daniel L Chen, Theodoros Evgeniou, and Anne-Laure Fayard. 2020. 'A Better Way to Onboard AI'. *Harvard Business Review*, 98, 4, 56–65.

Mahdi Bakhshayesh and Hamidreza Abbasianjahromi. 2023. 'Integrating text mining and system dynamics to evaluate financial risks of construction contracts'. *Artificial Intelligence and Law*.

Dovilė Barysė and Roee Sarel. 2024. 'Algorithms in the court: does it matter which part of the judicial decision-making is automated?' *Artificial Intelligence and Law*, 32, 1, 117–146.

Trevor Bench-Capon. 2022. 'Thirty years of Artificial Intelligence and Law: Editor's Introduction'. *Artificial Intelligence and Law*, 30, 4, 475–479.

[33] Note there are many commercial products based on, or inspired by, AI & Law research. It is, however, unclear exactly what kinds of techniques are being used in these systems, and honest evaluations of these systems are lacking.

Trevor Bench-Capon et al.. 2012. 'A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law'. *Artificial Intelligence and Law*, 20, 3, 215–319.

Floris Bex. 2024. 'AI, law and beyond: A transdisciplinary ecosystem for the future of AI Law'. *Artificial Intelligence and Law*.

Floris Bex and Henry Prakken. 2021. 'On the relevance of algorithmic decision predictors for judicial decision making'. In: *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*. ACM, São Paulo Brazil, 175–179.

Andrew Blair-Stanek, Nils Holzenberger, and Benjamin Van Durme. 2023. 'Can GPT-3 Perform Statutory Reasoning?' In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ACM, Braga Portugal, 22–31.

AnneMarie Borg and Floris Bex. 2021. 'Explaining Arguments at the Dutch National Police'. In: *AI Approaches to the Complexity of Legal Systems XI-XII* (Lecture Notes in Computer Science). Ed. by Víctor Rodríguez-Doncel, Monica Palmirani, Michał Araszkiewicz, Pompeu Casanovas, Ugo Pagallo, and Giovanni Sartor. Springer International Publishing, Cham, 183–197.

Daniel Braun. 2023. 'I beg to differ: how disagreement is handled in the annotation of legal machine learning data sets'. *Artificial Intelligence and Law*.

Stefanie Brüninghaus and Kevin D. Ashley. 2005. 'Generating legal arguments and predictions from case texts'. In: *Proceedings of the tenth international conference on Artificial intelligence and law*. ACM, Bologna Italy, 65–74.

Seth J Chandler, Emily A Muenster, and Daniel Lichtblau. 2023. 'Who really wrote Martin v. Hunter's Lessee? Applying a machine learning ensemble to uncover a Chief Justice's unethical behavior'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, 62–70.

Maxime C. Cohen, Samuel Dahan, Warut Khern-am-nuai, Hajime Shimao, and Jonathan Touboul. 2023. 'The use of AI in legal systems: determining independent contractor vs. employee status'. *Artificial Intelligence and Law*.

Jack G Conrad and John Zeleznikow. 2015. 'The role of evaluation in AI and law: an examination of its different forms in the AI and law journal'. In: *Proceedings of the 15th international conference on artificial intelligence and law*, 181–186.

Dennis Craandijk and Floris Bex. 2021. 'Deep learning for abstract argumentation semantics'. In: *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 1667–1673.

Dennis Craandijk and Floris Bex. 2022. 'Enforcement heuristics for argumentation with deep reinforcement learning'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. Issue: 5, 5573–5581.

Jet Cullen and Alan Bryman. 1988. 'The knowledge acquisition bottleneck: time for reassessment?' *Expert Systems*, 5, 3, 216–225.

Tonya Custis, Frank Schilder, Thomas Vacek, Gayle McElvain, and Hector Martinez Alonso. 2019. 'Westlaw edge AI features demo: KeyCite overruling risk, litigation analytics, and WestSearch plus'. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, 256–257.

Kristijonas Čyras, David Birch, Yike Guo, Francesca Toni, Rajvinder Dulay, Sally Turvey, Daniel Greenberg, and Tharindi Hapuarachchi. 2019. 'Explanations by arbitrated argumentative dispute'. *Expert Systems with Applications*, 127, 141–156.

Faraz Dadgostari, Mauricio Guim, Peter A Beling, Michael A Livermore, and Daniel N Rockmore. 2021. 'Modeling law search as prediction'. *Artificial Intelligence and Law*, 29, 3–34.

Jingpei Dan, Lanlin Xu, and Yuming Wang. 2023. 'Integrating legal event and context information for Chinese similar case analysis'. *Artificial Intelligence and Law*.

Robert Davison, Maris G Martinsons, and Ned Kock. 2004. 'Principles of canonical action research'. *Information systems journal*, 14, 1, 65–86.

Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2024. 'Ensemble methods for improving extractive summarization of legal case judgements'. *Artificial Intelligence and Law*, 32, 1, 231–289.

Fabiana Di Porto. 2023. 'Algorithmic disclosure rules'. *Artificial Intelligence and Law*, 31, 1, 13–51.

Laurence E Diver. 2021. *Digisprudence: code as law rebooted*. Edinburgh University Press.

Peter Fratrič, Mostafa Mohajeri Parizi, Giovanni Sileno, Tom van Engers, and Sander Klous. 2023. 'Do agents dream of abiding by the rules? Learning norms via behavioral exploration and sparse human supervision'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 81–90.

Leilei Gan, Kun Kuang, Yi Yang, and Fei Wu. 2021. 'Judgment prediction via injecting legal knowledge into neural networks'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35, 12866–12874.

Matthias Grabmair. 2017. 'Predicting trade secret case outcomes using argument schemes and learned quantitative value effect tradeoffs'. In: *Proceedings of the 16th Edition of the International Conference on Articial Intelligence and Law* (ICAIL '17). Association for Computing Machinery, London, United Kingdom, 89–98.

Neel Guha et al.. 2023. 'Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models'. *arXiv preprint arXiv:2308.11462*.

Zhilong Guo and Lewis Kennedy. 2023. 'Policing based on automatic facial recognition'. *Artificial Intelligence and Law*, 31, 2, 397–443.

Eliya Habba, Renana Keydar, Dan Bareket, and Gabriel Stanovsky. 2023. 'The Perfect Victim: Computational Analysis of Judicial Attitudes towards Victims of Sexual Violence'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 111–120.

Ivan Habernal, Daniel Faber, Nicola Recchia, Sebastian Bretthauer, Iryna Gurevych, Indra Spiecker genannt Döhmann, and Christoph Burchard. 2023. 'Mining legal arguments in court decisions'. *Artificial Intelligence and Law*.

Mireille Hildebrandt. 2018. 'Law as computation in the era of artificial legal intelligence: Speaking law to the power of statistics'. *University of Toronto Law Journal*, 68, supplement 1, 12–35.

Cong Jiang and Xiaolei Yang. 2023. 'Legal Syllogism Prompting: Teaching Large Language Models for Legal Judgment Prediction'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ACM, Braga Portugal, 417–421.

Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. 2024. 'Gpt-4 passes the bar exam'. *Philosophical Transactions of the Royal Society A*, 382, 2270, 20230254.

Daan Kolkman, Floris Bex, Nitin Naryan, and Manuella van der Put. 2023. 'Justitia ex machina: The impact of an AI system on legal decision-making and discretionary authority'. *Big Data & Society*, to appear.

Christie Lawrence, Isaac Cui, and Daniel Ho. 2023. 'The Bureaucratic Challenge to AI Governance: An Empirical Assessment of Implementation at U.S. Federal Agencies'. In: *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (AIES '23). Association for Computing Machinery, New York, NY, USA, 606–652.

Ludi van Leeuwen, Bart Verheij, Rineke Verbrugge, and Silja Renooij. 2023. 'Using Agent-Based Simulations to Evaluate Bayesian Networks for Criminal Scenarios'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 323–332.

Philip Leith. 2010. 'The rise and fall of the legal expert system'. *European Journal of Law and Technology*, 1, 1, 179–201.

Nicola Lettieri, Alfonso Guarino, Delfina Malandrino, and Rocco Zaccagnino. 2023. 'Knowledge mining and social dangerousness assessment in criminal justice: metaheuristic integration of machine learning and graph-based inference'. *Artificial Intelligence and Law*, 31, 4, 653–702.

Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. 2018. 'Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions'. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. Issue: 1.

Arpan Mandal, Kripabandhu Ghosh, Saptarshi Ghosh, and Sekhar Mandal. 2021. 'Unsupervised approaches for measuring textual similarity between legal court case reports'. *Artificial Intelligence and Law*, 29, 3, 417–451.

Gary Marcus and Ernest Davis. 2019. *Rebooting AI: Building artificial intelligence we can trust*. Vintage.

Marko Marković and Stevan Gostojić. 2023. 'Legal document assembly system for introducing law students with legal drafting'. *Artificial Intelligence and Law*, 31, 4, 829–863.

L Thorne McCarty. 1977. 'Reflections on" Taxman": An Experiment in Artificial Intelligence and Legal Reasoning'. *Harvard Law Review*, 837–893.

Scott McLachlan, Evangelia Kyrimi, Kudakwashe Dube, Norman Fenton, and Lisa C Webley. 2023. 'Lawmaps: enabling legal AI development through visualisation of the implicit structure of legislation and lawyerly process'. *Artificial Intelligence and Law*, 31, 1, 169–194.

Jack Mumford, Katie Atkinson, and Trevor Bench-Capon. 2023. 'Combining a Legal Knowledge Model with Machine Learning for Reasoning with Legal Cases'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 167–176.

Aileen Nielsen, Stavroula Skylaki, Milda Norkute, and Alexander Stremitzer. 2023. 'Effects of XAI on Legal Process'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ACM, Braga Portugal, 442–446.

Esther Nieuwenhuizen, Albert Meijer, Floris Bex, and Stephan Grimmelikhuisen. 2023. 'Explanations increase citizen trust in police algorithmic recommender systems: Findings from two experimental tests'. *Under review*.

Daphne Odekerken, Floris Bex, AnneMarie Borg, and Bas Testerink. 2022. 'Approximating stability for applied argument-based inquiry'. *Intelligent Systems with Applications*, 16, 200110.

Cándido Otero, Dennis Craandijk, and Floris Bex. 2023. 'ORLA: learning explainable argumentation models'. In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*. Vol. 19. Issue: 1, 542–551.

Frank Pasquale and Glyn Cashwell. 2018. 'Prediction, persuasion, and the jurisprudence of behaviourism'. *University of Toronto Law Journal*, 68, supplement 1, 63–81.

Joeri Peters, Floris Bex, and Henry Prakken. 2023. 'Model-and data-agnostic justifications with A Fortiori Case-Based Argumentation'. In: *19th International Conference on Artificial Intelligence and Law*. ACM Press, 207–216.

Sebastiano Antonio Piccolo, Panagiota Katsikouli, Thomas Gammeltoft-Hansen, and Tijs Slaats. 2023. 'On predicting and explaining asylum adjudication'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 217–226.

Henry Prakken, Floris Bex, and Anne Ruth Mackor. 2020. 'Editors' review and introduction: Models of rational proof in criminal law'. *Topics in Cognitive Science*, 12, 4, 1053–1067.

Henry Prakken and Rosa Ratsma. 2022. 'A top-level model of case-based argumentation for explanation: Formalisation and experiments'. *Argument & Computation*, 13, 2, 159–194.

Henry Prakken and Giovanni Sartor. 2015. 'Law and logic: A review from an argumentation perspective'. *Artificial Intelligence*, 227, 214–245.

Dasha Pruss. 2023. 'Ghosting the Machine: Judicial Resistance to a Recidivism Risk Assessment Instrument'. In: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 312–323.

Júlia Riera, David Solans, Marzieh Karimi-Haghighi, Carlos Castillo, and Caterina Calsamiglia. 2023. 'Gender Disparities in Child Custody Sentencing in Spain: a Data Driven Analysis'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 237–246.

Edwina L Rissland, Kevin D Ashley, and L Karl Branting. 2005. 'Case-based reasoning and law'. *The Knowledge Engineering Review*, 20, 3, 293–298.

Tyss Santosh, Shanshan Xu, Oana Ichim, and Matthias Grabmair. 2022. 'Deconfounding Legal Judgment Prediction for European Court of Human Rights Cases Towards Better Alignment with Experts'. In: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 1120–1138.

Arfah Habib Saragih, Qaumy Reyhani, Milla Sepliana Setyowati, and Adang Hendrawan. 2023. 'The potential of an artificial intelligence (AI) application for the tax administration system's modernization: the case of Indonesia'. *Artificial Intelligence and Law*, 31, 3, 491–514.

Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, and Pascal Hitzler. 2021. 'Neuro-symbolic artificial intelligence'. *AI Communications*, 34, 3, 197–209.

Jaromir Savelka. 2023. 'Unlocking Practical Applications in Legal Domain: Evaluation of GPT for Zero-Shot Semantic Annotation of Legal Texts'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ACM, Braga Portugal, 447–451.

Miriam Schirmer, Isaac Misael Olguín Nolasco, Edoardo Mosca, Shanshan Xu, and Jürgen Pfeffer. 2023. 'Uncovering Trauma in Genocide Tribunals: An NLP Approach Using the Genocide Transcript Corpus'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 257–266.

Marijn Schraagen and Floris Bex. 2019. 'Extraction of Semantic Relations in Noisy User-Generated Law Enforcement Data'. In: *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*, 79–86.

Marijn Schraagen, Floris Bex, Nick Van De Luijtgaarden, and Daniël Prijs. 2022. 'Abstractive Summarization of Dutch Court Verdicts Using Sequence-to-sequence Models'. In: *Proceedings of the Natural Legal Language Processing Workshop 2022*, 76–87.

Marijn Schraagen, Matthieu Brinkhuis, and Floris Bex. 2017. 'Evaluation of Named Entity Recognition in Dutch online criminal complaints'. *Computational Linguistics in the Netherlands Journal*, 7, 3–16.

Marijn Schraagen, Bas Testerink, Daphne Odekerken, and Floris Bex. 2018. 'Argumentation-driven information extraction for online crime reports'. In: *CKIM 2018 International Workshop on Legal Data Analysis and Mining (LeDAM 2018)*.

Monika Simmler, Simone Brunner, Giulia Canova, and Kuno Schedler. 2023. 'Smart criminal justice: exploring the use of algorithms in the Swiss criminal justice system'. *Artificial Intelligence and Law*, 31, 2, 213–237.

Carlos Soares, Stephan Grimmelikhuijsen, and Albert Meijer. 2023. 'Screen-level bureaucrats in the age of algorithms. An ethnographic study of algorithmically supported public service workers in the Netherlands Police'. *Information Polity*, Preprint, Preprint.

Rachel E Stern, Benjamin L Liebman, Margaret E Roberts, and Alice Z Wang. 2020. 'Automating fairness? Artificial intelligence in the Chinese courts'. *Colum. J. Transnat'l L.*, 59, 515.

Harry Surden. 2019. 'Artificial intelligence and law: An overview'. *Georgia State University Law Review*, 35, 19–22.

Richard E Susskind. 1986. 'Expert systems in law: A jurisprudential approach to artificial intelligence and legal reasoning'. *The modern law review*, 49, 2, 168–194.

Mehmet B. Unver. 2023. 'Rebuilding 'ethics' to govern AI: How to re-set the boundaries for the legal sector?' In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*. ACM, Braga Portugal, 306–315.

Hannes Westermann and Karim Benyekhlef. 2023. 'JusticeBot: A Methodology for Building Augmented Intelligence Tools for Laypeople to Increase Access to Justice'. In: *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law* (ICAIL '23). Association for Computing Machinery, New York, NY, USA, 351–360.

Alice Witt, Anna Huggins, Guido Governatori, and Joshua Buckley. 2023. 'Encoding legislation: a methodology for enhancing technical validation, legal alignment and interdisciplinarity'. *Artificial Intelligence and Law*.

Gizem Yalcin, Erlis Themeli, Evert Stamhuis, Stefan Philipsen, and Stefano Puntoni. 2023. 'Perceptions of Justice By Algorithms'. *Artificial Intelligence and Law*, 31, 2, 269–292.

Linwu Zhong, Ziyi Zhong, Zinian Zhao, Siyuan Wang, Kevin D Ashley, and Matthias Grabmair. 2019. 'Automatic summarization of legal decisions using iterative masking of predictive sentences'. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, 163–172.

Xiang Zhou, Qi Liu, Yiquan Wu, Qiangchao Chen, and Kun Kuang. 2023. 'LK-IB: a hybrid framework with legal knowledge injection for compulsory measure prediction'. *Artificial Intelligence and Law*.

Malte Ziewitz. 2016. 'Governing algorithms: Myth, mess, and methods'. *Science, Technology, & Human Values*, 41, 1, 3–16.