

Katja de Vries

Introduction to the De lege Yearbook 2021: Law, AI and Digitalisation

1 AI and Digitalisation: what is it?

Artificial Intelligence (AI) is everywhere – but *what* is AI? Those working with AI often discard the notion as a vague buzzword and instead prefer speaking about Machine Learning (ML), which is the dominant set of techniques driving the current AI revolution. ML, in turn, has deep ties to good old-fashioned statistical methods. As a well-known joke states: *When you're fundraising, it's AI. When you're hiring, it's ML. When you're implementing, it's statistics.* No wonder that the definition of AI is a contested topic in the proposed AI Act.¹ In the initial version, proposed by the Commission on the 21st of April 2021, it says in Article 3(1):

‘artificial intelligence system’ (AI system) means software that is developed with one or more of the techniques and approaches listed in Annex I² and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with

¹ European Commission, Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts (COM(2021) 206 final), 21 April 2021.

² Annex I states: *Artificial Intelligence techniques and approaches referred to in Article 3, point 1:*

(a) Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;

Criticised for being a too broad and vague grouping of techniques, the Council proposed to reformulate this in the presidency compromise text³ (29th of November 2021) into a much narrower definition, which mainly comprises ML systems that transform input data (such as passenger data) into output data (such as “potential terrorist” and “ordinary traveller”) according to a rule (that is, a model) that is bottom-up inferred from training data (that is, existing data relating to confirmed terrorists and ordinary travellers, respectively):

‘artificial intelligence system’ (AI system) means a system that (i) receives machine and/or human-based data and inputs, (ii) infers how to achieve a given set of human-defined objectives using learning, reasoning or modelling implemented with the techniques and approaches listed in Annex I⁴, and (iii) generates outputs in the form of content (generative AI systems), predictions, recommendations or decisions, which influence the environments it interacts with;

The definition of “AI system” will no doubt be further contested and refined during the remainder of the legislative process. From a computer science perspective, AI is as vague a notion as it ever was. Yet, discussing AI has become unavoidable, precisely *because* it is currently crystallizing as a concept within legal, political, policy and public discourse and as an object for legal and ethical regulation. Law, like always, is pragmatic. It does not aim to pinpoint a metaphysical essence of an AI system. Instead, it defines AI in a way that fits the regulatory purpose. For example, in the aforementioned definitions of ‘AI system’ from the proposed AI Act, the last part reads that its outputs ‘influence the environments it interacts with’. Given that the AI Act aims to regulate AI systems that could adversely impact on society or individuals, and explicitly excludes

(b) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;

(c) Statistical approaches, Bayesian estimation, search and optimization methods.

³ Council of the European Union, Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts – Presidency compromise text, 29 November 2021, available at: <https://www.statewatch.org/media/2963/eu-council-ai-act-compromise-text-14278-21.pdf>.

⁴ Annex I (n. 2).

any general-purpose AI from its scope,⁵ this definition makes sense from a regulatory perspective. Yet, if one was to create a dictionary entry for “AI”, an exclusion of non-influential, general-purpose AI would be an unwarranted limitation of the term’s scope.

It is often a conglomerate of techniques, applications or situations that is the object of regulation, not merely AI in a narrow technical sense. As Bruno Debaenst (Chapter 1) poignantly argues, the impact of the current AI revolution, that is, the fourth industrial revolution, on law, cannot be understood in isolation from earlier industrial revolutions. While some legal challenges posed by the AI revolution are truly novel, other challenges overlap with those raised by the third industrial revolution, that is, the digital revolution of computer automation. To underline this regulatory intertwinement, this book is entitled *Law, AI and Digitalisation*: addressing regulatory digitalisation challenges is often a necessary step before advancing to challenges posed by AI as such. An example is intellectual property (IP) questions, raised by avatars mimicking human behaviour, looks or speech that are created using AI techniques, so-called *synthetic* avatars. Such avatars can, for example, be used as *vocational* avatars, that is, as digital stand-ins for human professionals, such as teachers⁶ or actors.⁷ Vocational avatars can be created by training them on footage from flesh and blood equivalents, which involves making copies of that footage for training purposes. Anyone who wants to create vocational avatars will, thus, have to find training materials and tackle the question of who owns the copyright over them. These are questions that are already highly relevant in today’s world where teaching has become increasingly digitised due to the pandemic. Can a university re-use a recorded lecture if a teacher does not consent? Who owns the copyright over digitised teaching materials produced at today’s universities – the university teacher or the employer? And if it is the university teacher who holds the copyright, does the university have a right of usage? As Marianne Rødvei Aagaard (Chapter 10) argues, in Sweden, teachers are normally copyright holders of their teaching materials, but the university could, in a limited set of situations, also have some usage rights. While Rødvei Aagaard

⁵ Article 52a and Recital 70a AI Act (n. 3).

⁶ Kristin Houser (2018). The world’s first digital teacher just debuted in New Zealand. Tomorrow’s teachers won’t all be flesh and blood, The Byte, online available at: <https://futurism.com/the-byte/digital-teacher-new-zealand-will>.

⁷ Mathilde Pavis (2021). Rebalancing our regulatory response to Deepfakes with performers’ rights. *Convergence*, 27(4), 974-998.

does not discuss the reuse of teaching material as training material for a robot teacher, the question of the permissibility of such reuse is likely to develop along the same lines as more conventional forms of reuse. Even if the conclusion might be different (for example, because re-streaming a recorded lecture was an intended and foreseeable re-use, while training an avatar-teacher was not), the analysis of the avatar situation would still make use of the same legal framework and concepts.

The reuse of digital teaching materials is a clear example of continuity between legal challenges raised by the third industrial revolution (computer automation and digitalization) and the fourth industrial revolution (AI). In contrast, certain AI systems, such as the high frequency trading algorithms used in financial markets (Magnus Strand, Annina H. Persson, and Malou Larsson Klevhill, Chapter 22) or the AI-based automated grading systems, discussed by Cecilia Magnusson Sjöberg and Rebecka Weegar (Chapter 16) clearly go beyond the legal challenges raised by computer automation and digitalisation. Automated decision-making (ADM) and risk assessment systems, which are discussed by several authors in this volume, can be based on ML techniques, and more traditional systems, based on digitisation and automation. The choice of topics in this volume is characterised by a legal pragmatism: it does not only address AI in a narrow sense, but it also tackles more classical digitalisation and computer automation questions in as far as they continue to be relevant in an AI context.

2 AI and Digitalisation: where is it?

While the question of *what* AI is might be contested, the question of *where* it is can be answered quite easily: it is almost everywhere. During the last few years, AI systems have become so ubiquitously applied that they hardly left any aspect of life untouched. AI is used in surveillance, healthcare, smart cars, decisions relating to social welfare benefits, financial investments, generation of texts and other media, grading, research, warfare, fraud detection, border control, etc., etc. The list could easily be extended to fill the rest of the page. With every field that AI touches upon, it also raises a particular set of legal questions. These questions are often very fundamental questions: when activities, which used to be a matter of human discernment – fighting a war, deciding how to invest money, making a medical diagnosis or driving a car – are delegated to AI, it forces us to revisit the ground rules of these activities.

The contributions in this volume represent the multiplicity, broadness and richness of the legal questions raised by AI. They are organised into three parts: (I) AI, digitalisation and law: foundational explorations, (II) challenges posed by AI and digitalisation to particular areas of law, and (III) AI and digitalisation in practice: legal perspectives.

3 AI, digitalisation and law: foundational explorations (Chapters 1–7)

The first part of this volume (Chapters 1–7) discusses legal foundational questions raised by AI. How to deal with a new phenomenon? In Stanley Kubrick's movie, *Space Odyssey 2001*, bewildered apes gather around a rectangular black monolith that lands on the prehistoric earth, and as their fear makes place for curiosity, they begin to touch and explore the surface of the unknown object. AI, like Kubrick's black monolith, is a new phenomenon that needs to be approached with explorative questions. How does the impact of AI compare to earlier technological revolutions? How should we relate to AI? How do we steer it? What data do we feed it? The chapters gathered in *Part I: AI and Law – foundational explorations* all breathe an explorative spirit that is similar to the one that dominates the beautiful opening sequence of *Space Odyssey 2001*. In Chapter 1, *Bruno Debaenst* moves between legal history and legal futurology. He looks at the current AI revolution (the fourth industrial revolution) and compares its potential impact on law with the first, second and third industrial revolution. Chapters 2, 3 and 4 look into ways to relate to AI and the legal status granted to it. In Chapter 2, *Annika Waern* uses her field of expertise, Human-Computer Interaction, to present the different types of relations that humans can have with AI; she argues that not all types of relations are equally desirable from a societal and normative perspective. In Chapter 3, *Bert Lehrberg* discusses the question if AI could be attributed some form of legal personhood. He connects this question to three settings: algorithmic contracting, autonomous functioning vehicles and artificial general intelligence. In Chapter 4, *Anni Carlsson* addresses the question of legal personality for AI from a completely different angle; she looks for legal guidance in fiction, as 'law is also present in an imaginary

society'.⁸ She shows how humanoid robots can challenge the legal dichotomy of legal property and personhood. In Chapter 5, *Stanley Greenstein, Panagiotis Papapetrou and Rami Mochaourab* discuss the difficulties in building human values into the design of AI when the promoting of one human value is often at the expense of another competing human value. After a theoretical discussion of value sensitive design, the authors show how three values derived from the field of data protection, namely explainability, privacy and accuracy, are used for the construction of an AI system that predicts the diagnosis of future patients and that is trained on medical patient data. Training data is an indispensable basic ingredient for the creation of any AI model. The last two Chapters (6–7) in the first part of this volume deal with the question of where to find (training) *data* to fuel AI innovation. In Chapter 6, *Katja de Vries* discusses the options for a researcher who wants to use personal data. Is it better to use personal data (which entails compliance with data protection legislation) or is it possible to use non-personal surrogates, such as synthetic data or data of deceased people? In Chapter 7, *Bengt Domeij* throws light on some newly proposed pieces of EU legislation, such as the Digital Markets Act and the Data Act. He dives into the question of how the EU proposes to facilitate business-to-business data sharing and under which exceptional circumstances businesses could be legally forced to share industrial data.

In thinking about specific legal challenges following from AI and digitalisation, one could, broadly speaking, take one of the following two approaches: either one departs from a particular *field of law* and discusses the various challenges posed to it, or one takes the opposite direction by departing from AI and digitalisation in practice, that is, a particular *tool or field of application*, and then studying which legal challenges are evoked. This division is the organisational principle that informs the second and third part of this volume. It should, however, be underlined that the division is not watertight. Most contributions in this volume look at one or more legal fields, as well as at a particular AI application. Yet, in most contributions, one of these perspectives is more dominant. The *second part* of this volume (Chapter 8–14) gathers contributions where the dominant perspective is an engagement with the challenges posed by AI to particular legal fields, such as constitutional law, intellectual property law or international humanitarian law, whereas the contribu-

⁸ Jaakko Husa, 'Comparative law, literature and imagination: Transplanting law into works of fiction' (2021). 28(3) Maastricht J Eur Comp Law 371, 383.

tions gathered in the *third part* (Chapters 15–22) are characterised by predominantly departing from AI and digitalisation in practice, that is, a particular tool or application.

4 Challenges posed by AI and digitalisation to particular fields of law (Chapters 8–14)

In Chapter 8, *Markku Suksi* provides a thorough discussion of the use of ADM tools in relation to Finnish *constitutional and administrative law*, notably in relation to central concepts such as legality and rule of law, and makes some interesting comparisons to the Swedish legal framework. In Chapter 9, *Inger Österdahl* looks at the possibilities for regulating lethal autonomous weapon systems in the context of *the law in war*, or *international humanitarian law*, notably discussing the framework provided by the United Nations (UN), the Swedish policy position, and the controversial role of human control over advanced weapon systems. The authors of the following two chapters (10–11) all engage with AI and digitisation through the lens of *intellectual property law*. In Chapter 10, *Marianne Rødvei Aagaard* addresses a question that is of great practical importance in a time where the pandemic has made teaching increasingly digital: namely, who owns the copyright over digitised teaching materials – the university teacher or the employer? In Chapter 11, *Silvia Carretta* lays out the problems associated with content moderation of copyright-infringing material on the internet and the controversial Article 17 of the Copyright Directive 2019/790, which establishes a platform liability which, in practice, seems to make the use of automated upload-filters unavoidable. In Chapter 12, *Mikael Hansson* takes a *labour law* perspective. Can an AI system be an employer and hence have employer liabilities? Hansson discusses this thought-provoking question by looking at two Swedish labour law cases and one case decided upon by the Court of Justice of the European Union (CJEU) and concludes that AI cannot be held liable in the way a human employer can. In Chapter 13, *Vladimir Bastidas* looks at the practice of algorithmic personalised price-discrimination from an *EU competition law* perspective. When personalised pricing entails that an undertaking, to a certain extent, may determine market conditions, could this result in the applicability of Article 102 Treaty of the Functioning of the European Union (TFEU), which prohibits abuse by an undertaking in a dominant position and sets limits for the exercise of

market power. Through a detailed analysis of case law, Bastidas raises several pivotal questions, for example, if the abuse in Article 102 TFEU only refers to behaviour directed towards other competitors or if it could also include the differential treatment of end-consumers. All in all, Bastidas concludes that, in theory, a wide interpretation of Article 102 TFEU may include cases on personalised prices, even though such an interpretation does not seem very likely given the existing case law. Part 2 of this volume is concluded by a contribution by *Mattias Dahlberg* that departs from the field of *tax law*. In Chapter 14, Dahlberg discusses the problem that the digital economy requires hardly any physical presence in the market state and that giant tech companies can get away with paying very little taxes by establishments in low-tax jurisdictions, which are not representative of the markets that they serve. New legislation on the taxation of multinational enterprises proposes to break away from traditional taxation principles and allow market states to tax the income generated in that state. However, to determine the income generated in a state could require extensive surveillance of how consumers use digital services and could result in infringements on consumer privacy.

5 AI and digitalisation in practice: legal perspectives (Chapters 15–22)

Part 3 opens with two chapters looking at the legal implications of using *AI in educational settings*. In Chapter 15, *Liane Colonna* discusses the use of AI-tools in higher education (HE), such as remote-based proctoring or predictive learning analytics, from the perspective of the proposed AI Act.⁹ Colonna questions how well the risk-based approach of the AI Act, which categorises AI systems as entailing unacceptable, high, limited or minimal risks, fits the reality of educational AI, where it will often be difficult to categorise the AI system because of the unforeseen uses that can emerge in complex and large HE institutions. Colonna also observes that a university, in many situations, is likely to be both the provider and the user of an AI system. Two other sources of concern are that the AI Act has a largely technocratic approach which puts a lot of responsibility on the risk self-assessment by the AI developer that seems to exclude a participatory discussion with universities, students and teachers, and that

⁹ AI Act (n. 1).

the AI Act is unlikely to sufficiently mitigate the potential abuse caused by biometric technologies like facial recognition in educational settings. In Chapter 16, *Cecilia Magnusson Sjöberg and Rebecka Weegar* look at AI-based automated grading systems, both from a computer science and a legal perspective. They describe a pilot project at Stockholm University where a compulsory one-page written assignment was graded using an AI-model. The model was trained on assignments that were graded by human graders. The AI-assigned grades were compared with grades given by two human graders. Sjöberg and Weeberg discuss how the combination of AI and human grading could potentially lead to enhanced equal treatment of students.

The next two Chapters look at the use of *AI in healthcare*. In Chapter 17, *Santa Slokenberga* looks at three types of EU regulatory responses to medical ML used in paediatric care: the General Data Protection Regulation,¹⁰ the proposed AI Act¹¹ and the Medical Device Regulation.¹² She argues that none of the surveyed legal instruments contribute to furthering the development and availability of the devices directly and thus, the EU misses a chance to contribute to reducing the therapeutic gap in paediatric medical care. In Chapter 18, *Charlotte Högberg and Stefan Larsson* ask what role transparency and explainability of AI could have in relation to patients' rights and information flows in Swedish health care. They argue that, as is often falsely argued, the highest quality of health care is not opposed to transparency and that, in fact, the best possible health care cannot be achieved without transparency.

In Chapter 19, *Katarina Fast* looks at the use of *AI in child-related social services*. She gives an overview of several AI tools that have been used by social services to identify children at high risk of maltreatment. Most of these tools have been adopted in a context of struggles related to increases in caseloads, funding cuts and staff shortages and ambitions to increase digitalisation. Due to legal, ethical and public trust problems,

¹⁰ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data and the repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance) OJ L 119, 4.5.2016, 1–88.

¹¹ AI Act (n. 1).

¹² Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and the repealing Council Directives 90/385/EEC and 93/42/EEC (Text with EEA relevance) OJ L 117, 5.5.2017, 1–175.

many of these tools have been discontinued. Fast, then, raises the question of what would be needed to make such tools compatible with relevant children's rights as laid down in the UN Convention of the Rights of the Child, and other legislations, such as the European Convention of Human Rights, the EU Charter of Fundamental Rights, General Data Protection Regulation 2016/679 and the proposed AI Act.

The following two Chapters look at *AI-based decision-making and governance in the public sector*. In Chapter 20, *Stefan Larsson and Jonas Ledendal* present a comprehensive overview of the policy positions taken by the government and public authorities in Sweden towards the use of AI in the public sector. They compare the Swedish position with policies and legislative initiatives at the EU level, as well as with research on the use of AI in the public sector. One of their conclusions is the importance of creating a more specific legal and policy framework that articulates how AI-based decision (support) systems can be used in the public sector in a way that is in accordance with principles of good administration and the values steering government employees in Sweden ("*statliga värdegrunden*") and that preserves trust in public administration. In Chapter 21, *Johan Eddebo and Anna-Sara Lind* provide a lively encounter between a philosopher of religion and a legal scholar, discussing the question of if and how information, which affects discourses and opinion formation in the public sphere in potentially undesirable ways, should be governed by automated means. Eddebo and Lind introduce the notion of double intransparency, in relation to the influence exercised through algorithms within the framework of digital communication platforms and contemporary media technologies: it pertains to both our inability to access the algorithms as such, as well as the difficulties in reproducing and examining their actual effects. They discuss several possible mitigation strategies, such as the ones in the AI Act, and call for a deliberation that is both multidisciplinary and inclusive of all parts of civil society to consolidate the foundations of liberal democracy in an age of algorithmic content moderation.

The final contribution to this volume looks at the use of *AI in algorithmic financial trading*. In Chapter 22, *Malou Larsson Klevhill, Annina H. Persson and Magnus Strand* present an overview of the extremely complex national Swedish and EU legislative framework, regulating algorithmic high-frequency trading. One of the regulatory gaps identified by Klevhill, Persson and Strand is the lack of a clear civil liability regime: while administrative and criminal law sanctions can be good tools to promote

legal compliance in the financial sector, the investor who suffers damages following a mistake in a financial trading algorithm is not helped by that.

6 Concluding thoughts: Law in a world organised by ML perceptions

As I have argued elsewhere,¹³ one pivotal characteristic of ML (and thus of AI) in comparison to other technologies is the fact that it often revolves around *pattern* discovery, or, to put it in more anthropomorphic terms, it is a *meaning-attributing* technology. It perceives something *as* something. In order not to crash, a smart car needs to see a tree as a tree, a pedestrian as a pedestrian, a traffic sign as a traffic sign; an autonomous weapon system needs to distinguish foe from friends and civilians from military combatants; a financial investment AI needs to recognise a profitable investment opportunity from a bad one; a medical diagnostic AI tool should not mistake a malignant tumour and a benign one; a predictive model used by social services should correctly identify which child is at risk of being abused and which is not. This is the fundamental difference between a more conventional one like a knife and an AI-tool: the knife can cut a loaf of bread, but it does not *perceive* the bread as something (healthy or unhealthy, old or fresh, etc.). AI-tools help humans to make sense of the world: to see patterns and sometimes to propose new patterns: for example, one can imagine a kitchen-AI that not only recognises meals and gives the underlying recipe, but also suggests other recipes that might result in even better meals. Making sense of the world used to be a capacity belonging exclusively to humans, or, at the very least, to living beings: ‘On an insignificant background of reality, imagination designs and embroiders novel patterns: a medley of memories, experiences, free fancies, absurdities and improvisations’.¹⁴ What is so revolutionary about AI in the current AI-revolution is that it has introduced a category of

¹³ Katja de Vries, Privacy, due process and the computational turn. A parable and a first analysis. In M. Hildebrandt & K. De Vries (eds.), *Privacy, Due Process and the Computational Turn. The Philosophy of Law Meets the Philosophy of Technology* (Routledge, 2013) 9–38.

¹⁴ August Strindberg, *The Dream Play* (transl. Edwin Björkman; Charles Scribner’s Sons, 1912), online available at: <https://www.gutenberg.org/files/45375/45375-h/45375-h.htm> In the Swedish original: ‘... på en obetydlig verklighetsgrund spinner inbillningen ut och väfver nya mönster: en blandning af minnen, upplevelser, fria påhitt, orimligheter

pattern-recognising *tools* to support, or replace, humans in making sense of the world.

How does law regulate AI? How does law capture AI inside existing and new legal frameworks and concepts? The contributions in this volume show that this cannot be answered univocally, and that there are many answers. With law, the devil is always in the details. Law never operates through big ethical abstractions. Instead, it takes AI tools as they come: intertwined with other techniques and technologies, ranging from very simple automated systems to highly autonomous complex AI, mixing old legal questions with novel ones, strongly situated in practice with their own particularities, etc. This volume is, thus, a tribute to law's ways of dealing with AI in all its diversity and situatedness.

och improvisationer'. August Strindberg, Ett Drömspel in *Kronbruden. Svaneovit. Drömspelet*. (Stockholm, Iduns, 1902) 12.