

# Guide on the use of Artificial Intelligence-based tools by lawyers and law firms in the EU 2022



The project is co-funded by  
the Justice Programme  
of the European Union



## **PUBLISHER**

Council of Bars and Law Societies of Europe  
Rue Joseph II, 40  
B-1000 Brussels (Belgium)  
[www.ccbe.eu](http://www.ccbe.eu)

European Lawyers Foundation  
Fluwelen Burgwal 58,  
2511 CJ The Hague (The Netherlands)  
[www.elf-fae.eu](http://www.elf-fae.eu)

## **AUTHOR**

Peter Homoki

## **CONTRIBUTORS**

Britta Kynast  
Christian Lemke  
Iain G. Mitchell QC  
Jiří Novak  
Carla Secchieri  
Thierry Wickers

### **European Lawyers Foundation (ELF)**

Alonso Hernandez-Pinzon Garcia  
Jonathan Goldsmith

### **Council of Bars and Law Societies of Europe (CCBE)**

Martin Sacleux  
Simone Cuomo

The CCBE and ELF hold the copyright for this report, without excluding the European Union's right in relation to article I.7 of the project's grant agreement.  
For more information, please address your query to [info@elf-fae.eu](mailto:info@elf-fae.eu)



# Table of Contents

|                                                                                                                                                 |    |
|-------------------------------------------------------------------------------------------------------------------------------------------------|----|
| List of abbreviations used .....                                                                                                                | 4  |
| Executive summary.....                                                                                                                          | 5  |
| 1. Introduction: Background and aims of the AI4Lawyers project.                                                                                 | 8  |
| 2. What opportunities do AI and related technologies offer to small law firms?.....                                                             | 10 |
| 2.1 AI and machine learning.....                                                                                                                | 10 |
| 2.2. Why is it important for society that lawyers make use of AI tools? .....                                                                   | 11 |
| 2.3 Other novel technologies besides AI.....                                                                                                    | 12 |
| 3. Explaining some basic terminology for AI tools: corpus, datasets, benchmarks and models, linguistic tools and knowledge representation ..... | 14 |
| 3.1 Dataset, training set, corpus and training methods.....                                                                                     | 14 |
| 3.2 Real-time training and end-user training capabilities .....                                                                                 | 15 |
| 3.3 Benchmarks and claims on the performance of a tool .....                                                                                    | 15 |
| 3.4 Why are linguistic (NLP) tools and knowledge representation important for lawyers? .....                                                    | 16 |
| 3.4.1. The role of NLP and linguistic tools in AI tools for lawyers .....                                                                       | 16 |
| 3.4.2. Legal knowledge representation beyond the surface level of the text .....                                                                | 17 |
| 4. Opportunities for bars and law societies .....                                                                                               | 19 |

|                                                                                                                                      |           |
|--------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b>5. Categories of AI tool features and connecting it to lawyers' typical activities</b>                                            | <b>20</b> |
| 5.1. Drafting support tools                                                                                                          | 20        |
| 5.1.1. Writing assistance tools                                                                                                      | 20        |
| 5.1.2. Document assembly tools                                                                                                       | 21        |
| 5.1.3. Tools for turning legal data and knowledge bases into text                                                                    | 23        |
| 5.2. Document analysis                                                                                                               | 24        |
| 5.2.1. Introduction                                                                                                                  | 24        |
| 5.2.2. Understanding by way of classifying text or parts                                                                             | 25        |
| 5.2.3. Analysis based on information extraction: extracting time, relations (citations), content of contractual provisions and facts | 26        |
| 5.2.4. Combining classification and extraction for document understanding and analysis purposes                                      | 26        |
| 5.3. Text retrieval and analysis of case law and legislation                                                                         | 27        |
| 5.3.1. Introduction                                                                                                                  | 27        |
| 5.3.2. On three different objectives of legal analytics                                                                              | 29        |
| 5.3.3. Advanced searching techniques beyond the text: semantic search and argument mining                                            | 31        |
| 5.3.4. Analysis of activities of participants based on case law                                                                      | 32        |
| 5.4. Speech-to-text tools                                                                                                            | 33        |
| 5.5. Chatbots                                                                                                                        | 34        |
| 5.6. Assistance in internal office administration by AI tools                                                                        | 36        |
| <b>6. Scenarios</b>                                                                                                                  | <b>38</b> |
| 6.1. Introduction                                                                                                                    | 38        |
| 6.2. Bilateral contract negotiations on a platform and recording time                                                                | 38        |
| 6.3. Client meeting and intake                                                                                                       | 39        |
| 6.4. A quick lease agreement is needed                                                                                               | 40        |
| 6.5. Preparing for court work                                                                                                        | 41        |
| 6.6. A service for a professional entrepreneur                                                                                       | 41        |
| 6.7. Court work for Mr Beta and submission of the package to the court                                                               | 42        |

|        |                                                                                                                                     |    |
|--------|-------------------------------------------------------------------------------------------------------------------------------------|----|
| 7.     | Risks to professional obligations when using AI tools . . . . .                                                                     | 43 |
| 7.1.   | Introduction. . . . .                                                                                                               | 43 |
| 7.2.   | Risks of technological nature . . . . .                                                                                             | 44 |
| 7.2.1. | Risks arising from using cloud computing and online platforms to provide access to AI tools . . . . .                               | 44 |
| 7.2.2. | Relying on results without proper explanation and understanding, and others risks relevant to the performance of AI tools . . . . . | 45 |
| 7.2.3. | Risks to privacy . . . . .                                                                                                          | 47 |
| 7.3.   | Risks arising from professional obligations . . . . .                                                                               | 49 |
| 7.3.1. | Risks to professional competence: dangers of trying out new technology . . . . .                                                    | 49 |
| 7.3.2. | Risks to professional competence: integrating technical and human processes, balancing promises with actual capabilities . . . . .  | 50 |
| 7.3.3. | Risks related to professional secrecy obligations of the lawyer . . . . .                                                           | 51 |
| 7.3.4. | Risks related to the independence of the lawyer. . . . .                                                                            | 51 |
| 8.     | Conclusions . . . . .                                                                                                               | 53 |
| 9.     | Bibliography . . . . .                                                                                                              | 54 |



## List of abbreviations used

|      |                                                                                                                                     |
|------|-------------------------------------------------------------------------------------------------------------------------------------|
| AI   | Artificial intelligence                                                                                                             |
| API  | Application programming interface                                                                                                   |
| CCBE | Council of Bars and Law Societies of Europe                                                                                         |
| dApp | A decentralised application running on a specific DLT-based network (blockchain)                                                    |
| DLT  | Distributed ledger technology (including blockchain, see footnote 17)                                                               |
| ELF  | European Lawyers Foundation                                                                                                         |
| FMCG | fast-moving consumer goods (see footnote 73)                                                                                        |
| NFC  | near-field communication (see footnote 74)                                                                                          |
| NFT  | non-fungible token (a type of blockchain unit that is not interchangeable the same way as e.g. cryptocurrencies on blockchains are) |
| NLP  | natural language processing                                                                                                         |
| OCR  | optical character recognition (an automated way of converting images of text to computer readable text)                             |



## Executive summary

**This guide aims to provide information on how lawyers will be able to use the opportunities provided by AI tools and how such tools could help the business processes of small firms. Its objective is to provide lawyers with some background to understand what they can and cannot realistically expect from these products. This guide aims to give a reference point for small law practices in the EU against which they can evaluate those classes of AI applications that are probably the most relevant for them.**

This guide is the third in the AI4Lawyers project, following the “Overview on average state of the art IT capabilities and comparison with best practices United Kingdom, USA and Canada” and the deliverable of the second phase, “Report on the barriers and opportunities in the use of natural language processing tools in small legal offices”. The approach of the present guide is to be more didactic than product-centred, providing practical information as to which particular tools to use and how.

These tools provide opportunities for smaller law firms to be empowered to respond to a more digitised society. New ways of automation can provide new chances for lawyers to improve their workflows. However, there are considerable difficulties: small law firms may lack sufficient IT budget and access to consultants, and this problem is exacerbated by the peculiarities of the fragmented legaltech market in the EU.

To understand how these tools can be useful or harmful, some basic terminology has first to be understood by lawyers who are considering the use of such tools. For that purpose, this guide gives a short, non-technical explanation for expressions like corpus, datasets, benchmarks and models. This should help lawyers to evaluate claims made by publishers, for instance when referring to the excellent results of a product against a specific benchmark, or when pointing out why research in linguistic or natural language processing tools is the cornerstone of such products.

After reviewing these foundations, **the guide gives an overview of the AI applications in six different categories that are probably of the greatest relevance to lawyers.**

The first category relates to **drafting support tools**, itself divided into writing assistance tools, document assembly tools and a more generic division of text generation from non-text data.

**Writing assistance tools** (including those recommending a specific style or providing access to reusable text snippets, checking of citations etc.) mostly are integrated into word processors, which are currently still the main user interface for lawyers.

**Document assembly software** is designed to facilitate the automation of the construction of documents from template texts based on specific conditions. Such assembly first requires proper authoring of templates with a business logic setting out which texts to use at which times, which questions to ask etc. Of course, lawyers are interested in reusing clauses as often as possible in as many cases as are feasible, because otherwise automation cannot achieve economies of scale. However, such a general approach requires the defining of text templates in a very general way. This results in a much greater complexity during authoring, and makes it more difficult to train lawyers themselves to use these tools. The market for document assembly tools is a relatively mature market, available to most lawyers since

the 1990s, but its use is still not widespread, probably due to complexities in authoring and the lack of language specific features (at least outside English).

**Generating text from non-text data into legal texts** is not extensively used by small firms in the EU at this time. Such use could include new tools for text generation, such as turning data recorded in case management software concerning, for example, evidence, into natural text or submissions, or to incorporate arguments made by a party into a structure based on which responses can be submitted to the court.

Under the category of **document analysis**, we seek to explain how AI tools attempt to understand and analyse documents submitted by the lawyer, including how documents and their parts are classified and how important provisions are extracted, and how these tools are designed in a way to make these techniques work together and deliver an automated analysis of texts which have been submitted (such as summarising texts for due diligence purposes, highlighting problematic provisions or even making suggestions for contracts reviewed).

Next, we discuss **text retrieval, case law analysis and legislation analysis tools**. The first part of this discussion explains how and why text retrieval applied to legal text searches which had originally provided full text keyword-based searches had to meet expectations that only the most relevant results would be provided, and how this changed the objective of text retrieval software. Thanks to advances in natural language processing, text retrieval software has been enhanced to enable it to provide lawyers with access to information in case law and legislation previously impossible to comprehend due to its mass, making new connections possible between texts. The process of focusing on quantitative data buried deep in case law brought about the introduction of a new class of tools called “analytical justice tools”, where the search is no longer focused on retrieving relevant textual information, but on amounts of money (in claims or damages calculations), lengths of prison sentence and the extent of other penalties. Quantitative search tools themselves have made yet another class of tools possible in case law analysis: predictive tools, which, based on case law, attempt to give an estimate of various quantitative results of the case. After discussing these three levels of legal analytic tools, the chapter ends with a presentation on how a semantic search may lead to argumentation-based searches and what the advantages of argument mining tools are for lawyers.

Even today, **speech-to-text tools** are already available for lawyers in all of the official languages of the EU, but that does not mean that there is no room for improvement within this category, such as in relation to accurate recognition of speech without prior training in that user’s voice, or in transcribing recordings made in noisy environments, or of multiple speakers, or even in automating the creation of minutes of meetings from recordings.

With the enormous growth in the use of social media and messaging applications, interest in and use of **chatbots** have also increased, with ample potential for small law firms as well. To enable lawyers to become more familiar with these tools, the guide explains how these tools usually work and interact with users (including understanding natural text in chats), and what is expected from those who implement these tools (like defining intents and entity types, and also reactions to any of those intents or entities so identified). Of course, the privacy and confidentiality risks of such tools require to be addressed as well.

The final category is focused on those AI tools which can assist lawyers in making some of their internal administrative processes more streamlined and efficient. The first subclass of this category relates to **automated time tracking of lawyers**, which is not only useful for efficiency purposes, but might also at the same time pose dangerous risks to privacy. It is imperative for law firms to understand what these tools actually capture as they operate, to inform their employees or other lawyers of such tracking, and to comply with all the relevant privacy protection requirements. Another subclass of applications relates to **simplifying and standardising data streams acquired from time tracking** before such data are presented to the clients for billing, and tools that help lawyers in filing documents with appropriate metadata without having to spend too much time on such clerical tasks.

After in-depth discussion of the various categories of these applications, **six future scenarios** are presented in a narrative form. These scenarios are taken from the imaginary life of an ideal lawyer from the future who uses as many of the aforementioned tools as feasible, in mostly realistic scenarios.

In the last chapter, **seven risk areas** are highlighted of which lawyers should be aware, partly technical risks and partly risks of breaching lawyers' professional obligations (building upon core principles from the Charter of Core Principles of the European Legal Profession and Code of Conduct for European Lawyers). As long as there is an expectation that lawyers will increase their processing capabilities, for example by the use of new technical tools, there will be a requirement for lawyers to do this in a way that respects the core principles of the profession, such as the duty of client confidentiality, expectations of competence and the independence of the profession.

One of the biggest risks in using these new tools is related to the manner in which such tools are generally delivered, **cloud computing**, largely because obtaining data from the cloud is much harder than putting the data in, and the more successful a cloud provider becomes for lawyers, with integrated offerings, the stronger such service providers become when facing lawyers.

Another relevant technical risk of AI tools is the **lack of transparency and explainability**, a problem which affects most AI users, but which has special relevance for lawyers. The application of the rule of law brings a requirement for a strong-sense interpretability of the results of a model, which can be difficult to achieve, and the word embedding language model techniques (foundational models) further exacerbate this problem which has no solution at this time. There are also other, technical risks posed by these applications, including the **brittleness of training data** (severely degraded performance in situations not encountered during the training) and **discrimination** (due to the training data used, the architecture chosen or simply due to unexpected changes in the use of the tool, such as in cross-jurisdictional uses).

The **privacy risks** from AI tools are partly the same as those with regard to cloud computing (such as the possibility of access by law enforcement agencies or other actors), but there are also AI specific risks through service providers reusing data at their disposal in ways that are not fully transparent to the lawyer. The dangers of problematic reuse include the re-identification of a supposedly anonymous dataset, or the technical risk that, from the results of a trained model, sensitive inferences can be made in relation to the original training data as well.

With regard to professional obligations, the first core principle examined is the **competence of a lawyer**. Like many other new technologies, this affects lawyers both externally and internally. Lawyers are often expected to assume how society and courts will react to new technical possibilities, to new tools available, even when there is simply not enough data make competent guesses about the social impact of such tools. Further, due to a fear of missing out, lawyers may be tempted to start using new tools that are not yet properly tested in the given jurisdiction, which may lead to unexpected data breaches etc. Although effective AI tools can multiply the capabilities or reach of a law firm, the obligation of professional competence requires lawyers not to overcommit or overpromote themselves: a traditional law firm model still requires that lawyers have a human understanding of what clients need, which is often not the same as what clients may say they need. This raises differences from the usual business model of a legaltech provider of consumer services, and clients should be made aware of such differences when a legal service provider provides them with services.

**Various risks to client confidentiality** have been pointed out in various parts of the guide, including in the in-depth discussion of specific categories of tools (when using chatbots, document analytics or speech-to-text tools), but also in relation to cloud computing and the privacy risks of tools. But client confidentiality is not only a technical matter, it is also a strategic consideration for lawyers when they choose their tools, their IT architecture and the clients they serve. Advances in security in recent decades have provided ample evidence that, for the most sensitive situations, lawyers have to be aware that some clients can be better protected in an environment which is as offline as possible. Even if access by third-parties can never be ruled out, offline access can be made a lot more challenging for third-parties, at least compared to any promises of state-of-the-art data centres or SOC2 reports and certifications.



# 1. Introduction: Background and aims of the AI4Lawyers project

In 2018, the European Commission launched a European initiative on Artificial Intelligence (AI) <sup>1</sup> having the aim of preparing for socio-economic changes brought about by AI by, among other things, encouraging the modernisation of education and, in anticipation of changes in the labour market, supporting labour market transitions.<sup>2</sup>

The Council of Bars and Law Societies of Europe (CCBE), which represents the bars and law societies of its 32 member countries and 13 further associate and observer countries, and through them more than 1 million European lawyers, has since 2014 had a number of special committees and working groups dedicated to following the impact of new technologies, and, since 2016, has published several studies on the effects of AI. After having adopted the *CCBE Considerations on the Legal Aspects of Artificial Intelligence* in 2020 (“CCBE Considerations”),<sup>3</sup> the CCBE has submitted a number of project proposals to the European Commission (EC) to carry out further in-depth studies.

The Council of the European Union adopted in 2019 the 2019-2023 Action Plan on European e-Justice, which sets out a list of projects and initiatives (‘actions’) to be implemented as part of the 2019-2023 European e-Justice Strategy. The Action Plan also indicates the goals of individual actions and the envisaged activities, the participants, and the expected contributions of the stakeholders involved (citizens, companies, legal practitioners and judicial authorities). The drafting of a guide on the use of AI by lawyers in the EU was mentioned in the Action Plan under the possible actions to be implemented under ‘Artificial Intelligence for Justice’. Taking this into account, and based on the call for proposals for action grants 2019, the CCBE and the European Lawyers Foundation (ELF) submitted a project proposal on Artificial Intelligence for Lawyers (AI4Lawyers).<sup>4</sup> The CCBE and ELF were awarded an EU Grant to implement that project, which started in 2020.

The present guide is a deliverable of the AI4Lawyers project. The first phase of the project resulted in a study entitled “*Overview on average state of the art IT capabilities and comparison with best practices United Kingdom, USA and Canada*”<sup>5</sup> and the second phase ended with the “*Report on the barriers and opportunities in the use of natural language processing tools in small legal offices*”.<sup>6</sup>

This guide is intended to contain practical information for lawyers and small sized law practices on how to integrate AI applications into their everyday work. However, in this case, practical information has a special meaning. The main goal is to assist in understanding how some currently popular categories of such tools work and how they can be put at the service of lawyers in a way that does not undermine their professional obligations.

<sup>1</sup> See the list of abbreviations on page 5.

<sup>2</sup> European Commission, ‘[Communication from the Commission: Artificial Intelligence for Europe](#)’ (2018) accessed 19 November 2021.

<sup>3</sup> Council of Bars and Law Societies of Europe, ‘[CCBE Considerations on the Legal Aspects of AI](#)’ (2020) accessed 19 November 2021.

<sup>4</sup> This short title has been used for the project since May 2019. It is not to be confused with e.g., the project of the University of Oxford, called AI4LAW (used at least since around the same time). The projects are not related.

<sup>5</sup> Péter Homoki, ‘[Overview on Average State of the Art IT Capabilities and Comparison with Best Practices United Kingdom, USA and Canada](#)’ (Council of European Bars and Law Societies (CCBE), European Lawyers Foundation).

<sup>6</sup> Pál Vadász and others, ‘[A Report on the Barriers and Opportunities in the Use of Natural Language Processing Tools in Small Legal Offices](#)’ (Council of European Bars and Law Societies, European Lawyers Foundation).

Therefore, the approach of this guide is intended to be more didactic than product-centred, and so no recommendations are given as to which particular tools to use and how to use them. Given the immense differences within EU member states,<sup>7</sup> such an approach would be unfeasible and short-lived. If we take a look at the products currently available in the various markets for small law firms in the EU, most of the solutions are not a good indication of what might be achieved in the coming years.

The aim of this guide is to provide information on how lawyers will be able to use the opportunities provided by these new technologies, and what kind of AI tools could help the business processes of small firms. We have the aim of making lawyers more informed as to the nature of present and future products and providing some reference points for the evaluation of AI products built on the AI/NLP (natural language processing) applications mentioned in this guide.

The guide starts by dedicating a few pages to explaining some basic technical terms in sections 2-3, and also to discussing why the use of AI tools by lawyers is important not only for lawyers themselves, but also for the rest of society. To introduce the more technical parts, we first provide detailed explanations for certain categories of tools in section 5, and later present some imaginary scenarios providing a narrative around how we as lawyers may in the future use the tools (see section 6). The guide concludes by setting out the professional risks of which lawyers should be aware (see also section 6).

---

<sup>7</sup> Homoki (n 5) 10.



## 2. What opportunities do AI and related technologies offer to small law firms?

This part first gives a short explanation about why AI can mean so many things, and then explains the use of NLP/AI in specific contexts, automation techniques not yet used in the legal field ( “innovative uses”), regardless of whether the tools are rule-based tools, or based on machine learning techniques from 1986 or on the latest deep-learning models of today.

### 2.1 AI and machine learning

Artificial intelligence as a term does not refer to specific techniques but rather to a general objective of tools used for the carrying out in an automated way activities previously undertaken by humans.<sup>8</sup> This is a very wide field of research which has been ongoing for many decades in many different branches of science, including in computer science or in biology (under many different names and objectives). One area of typical application of AI that is relevant to the work of lawyers is referred to as natural language processing (abbreviated to NLP).<sup>9</sup> This AI term is also used to target potential customers interested in enhancing their current automation capabilities with specific software (in other words, as a marketing term), and this is the use with which most lawyers will be familiar. In this sense, AI is a term that encompasses all the various techniques and products used for replacing or enhancing certain human capabilities, such as automated navigation of a vehicle in an area where obstacles have to be identified in the same way as humans perform the task, or processing the content of a document written in a natural language, etc.

*AI is a term that encompasses all the various techniques and products used for replacing or enhancing human capabilities*

In recent years, lawyers’ attention has been repeatedly drawn to tools that use techniques called machine learning. Machine learning and AI are not the same and they should not be used interchangeably, nor are they so used in this guide. In order to elucidate the relationship between how we used AI and machine learning in this guide, some explanation is needed.

“Machine learning” is itself also a composite term which encompasses any procedure (including algorithms) designed so that the procedure changes itself based on the values of a set of examples (this set is called the training data, training set or the dataset). In other words, certain parameters of the procedure change, based on the examples provided during the training activity, and these adapted parameters are used for future runs of the procedure. The procedure is thus expected to provide better

<sup>8</sup> Council of Bars and Law Societies of Europe, ‘CCBE Considerations on the Legal Aspects of AI’ (n 3) 8–10.

<sup>9</sup> For further details on NLP, see *ibid* 8–17.

results than would have been achieved by the procedure without training. The results of the training (which could be e.g. a simple decision tree<sup>10</sup>) is often called the “trained model”.

Many fields of current AI do not rely on such techniques, meaning that they do not use machine learning, but are nevertheless able to perform their task of imitating some human capabilities well, such as acting on the content of an unknown document (“understanding the document”). For example, argument mining (see section 5.3.3) is viewed as a field of research in AI, and in specific cases it is possible to identify most logical connectives within sentences based on a list of text patterns predefined by humans (e.g. regular expressions). This is often called a rule-based approach, which is contrasted with a machine learning-based approach: it is not machine learning, but AI. Also, machine learning techniques (procedures) can be used for any purpose, without any intention to replace or enhance human capabilities.

One has to be aware that, just as in many fields of law, there is no single truth as to what is the best definition of AI, nor as to its relationship with machine learning. However, from the viewpoint of this guide and that of lawyers, we believe that the best approach is to define and delineate the two terms as above.

The most visible breakthroughs in recent years in different AI applications have all been related to machine learning, more specifically, to neural network-based machine learning techniques.<sup>11</sup>

## 2.2. Why is it important for society that lawyers make use of AI tools?

In the introduction to his book *New Laws of Robotics*, Frank Pasquale said that “*one way to alleviate the democratic crisis of expertise—the tension between aloof technocrats and passionate populists—is to empower local professionals*”.<sup>12</sup> In the report of phase one of the project AI4Lawyers,<sup>13</sup> ample evidence has been provided in the form of statistics that smaller legal practices have a very important role in the task of supporting the rule of law.

AI tools are a great opportunity to empower law firms of all sizes to be able to respond to changing client requirements (including changing preferred channels of communication) and the increase in the amount of data generated at the level of society, including the ever-increasing amount and diversity of digital evidence which is created and which requires to be processed. From one perspective, this is an opportunity for a number of smaller firms to compete successfully with larger firms in sectors that they previously could simply not have served due to technical bottlenecks and lack of capacity. Based on more extensive automation, IT tools could provide a new chance for lawyers to improve further their workflows, to make their deliverables more consistent, and to increase the added value of their work.

A human expert with a human understanding of the full context of a client is able to provide a much higher quality of service to clients, compared with automated assistance in a restricted set of piecemeal activities that can be fully automated by software, but where clients have to rely on their own judgement and self-help.

10 A representation of possible choices usually visualised in a tree-like format with branches dividing into smaller branches and ending in “leaves”, like “if the text contains the word “service”, then the result should be “a service contract”, otherwise “work contract”, where the decision on the text containing the word “service” is a branch with two leaves, being a minimalistic tree.

11 Vadász and others (n 6).

12 Frank Pasquale, *New Laws of Robotics: Defending Human Expertise in the Age of AI* (The Belknap Press of Harvard University Press 2020).

13 Homoki (n 5) 8.

*A human expert with a human understanding of the full context of a client is able to provide a much higher quality of service to clients, compared with automated assistance in a restricted set of piecemeal activities that can be fully automated by software, but where clients have to rely on their own judgement and self-help.*

At the same time, small law firms face considerable challenges due to not having an adequate IT budget and not having access to consultants. One of the gaps identified in the report of the first phase highlights the dangers which arise from those small law firms in the EU failing to use most categories of AI tools and emphasises the important difference between the markets for such tools in the US, the UK and the countries of the EU. In England and in the states of the USA, there is already a potentially healthy market for many AI-assisted tools used by larger firms, which also puts smaller law firms in a better position to access such tools if they want to.<sup>14</sup> Nevertheless, due to legal and linguistic differences, most of these products will not be marketed for law firms in many of the Member States of the EU. These barriers were also confirmed by the phase 2 report.<sup>15</sup>

For the reasons stated above, it is not only in lawyers' best interest to understand how such tools work or in what directions they may develop in the future, but it is also important for society at large to enable as many lawyers as possible to use such tools effectively in the interest of their clients.

## 2.3 Other novel technologies besides AI

When discussing AI-based tools, this guide cannot avoid discussing the effects of cloud computing technology on the work of lawyers. Cloud computing as a technology of delivery (of making new tools available to lawyers) is inextricably intertwined with the use of AI/NLP tools. The simplicity of accessing a new service without the end-users having to undertake a great deal of configuration offers a tremendous advantage for non-technical users like sole practitioners and smaller law firms, neither of which tend to have adequate access to IT professionals and consultants. This makes cloud computing a very attractive solution for this segment of law firms,<sup>16</sup> until the disadvantages of cloud computing start to outweigh the advantages due to the increasing size of the firm. Some of these risks are discussed in section 7.

There are also many other novel technologies that fall outside the issue of cloud computing. Among the technologies that imply long-term changes at lawyers' workplaces, are included the Internet-of-Things, distributed ledger/blockchain technologies<sup>17</sup> and also, the many changes to be expected in the world of electronic identity and bring-your-own-identity solutions (see the new proposed eIDAS amendments<sup>18</sup>).

<sup>14</sup> *ibid* 47.

<sup>15</sup> Vadász and others (n 6) 20–23.

<sup>16</sup> Homoki (n 5) 9–10.

<sup>17</sup> Internet-of-Things means a computer network of physical objects equipped with built-in technologies for interacting with each other or with their environment, such as smart home devices, voice assistants, activity trackers etc. We use the terminology of DLT and blockchain as defined in ISO 22739:2020. Many authors use different definitions and make incompatible distinctions between what they see as the “correct” term. In this simplified sense, a *ledger* is an information store that keeps records of transactions that are intended to be final, definitive and immutable, and a *distributed ledger* is a ledger shared across a set of computers (DLT nodes containing a copy of the ledger each) and synchronized between these nodes using a consensus mechanism (a way for the majority of DLT nodes to validate a transaction), and a *blockchain* is a type of distributed ledger which is organised in an append-only sequential chain using cryptographic links between each blocks. The best-known crypto-currencies are based on such blockchains technologies, they are a type of blockchain applications. DLT is a more generic technique that can be implemented in many different ways – but these are usually not as interesting or visible to the public.

<sup>18</sup> ‘[Proposal for a Regulation of the European Parliament and of the Council Amending Regulation \(EU\) No 910/2014 as Regards Establishing a Framework for a European Digital Identity](#)’ accessed 21 November 2021.

The in-depth analysis of such tools would warrant separate studies, but currently, these novel technologies seem to affect the specific workflows of lawyers very differently compared to the much broader concept of AI as defined above.

For instance, the Internet-of-Things as a technology will change the operations of a law firm, but what we can currently see in this field is not specific to the legal profession. Profession-specific changes will probably be a consequence of changes which are yet to come at the courts and other authorities, such as new types of digital evidence becoming available or a new sensor embedded within the workflow of courts and other authorities. These products will definitely provide new opportunities in gathering or processing evidence for existing legal activities like court or administrative procedures, but we are not yet in a position to see whether such processes will be capable or not of replacing considerable parts of existing business processes at law firms.

Similarly, the advantages of using distributed ledger (DLT) technology is not specific to the operation of law firms. Law firms are private operators of their own infrastructure, as a result of which the distributed operation of their own infrastructure is rarely seen as desirable. But once a public registry, such as a land registry, turns into a DLT-based registry, it will become a matter of training the lawyer as to how to initiate the required changes in the registry (provided lawyers are entitled to do so in a given jurisdiction), how to ask for judicial review, how the DLT registry will enable those reviews and enforce third-party decisions.

Many clients expect lawyers working for individuals to understand how the most popular crypto-assets infrastructure works in practice, because the lawyer is expected to give advice on such everyday matters as inheritance or divorce where a crypto-asset is involved. The lawyers are expected to give advice on how to bridge the gap between law and practice, regardless of whether that is technically possible. But these technical questions are all related to a specific type of application on a DLT, the design and implementation of which is not a question of strategic guidance for the operation of the law firm using these technologies.

As soon as a specific decentralised application of a DLT becomes socially important, and is used by a large number of people, e.g. consumers, banks, companies etc., people start to have a need to enforce existing rules and regulations in relation to these technologies as well. One may store a crypto asset anonymously, but as soon as people want to buy a house or a car or jewellery with the value of the asset, it will become subject to the same anti-money laundering rules and taxation etc. as with any traditional asset.

*As soon as a specific decentralised application on a digital ledger technology becomes socially important, and is used by a large number of people, e.g. consumers, banks, companies etc., people start to have a need to enforce existing rules and regulations on these technologies as well.*

In summary, the effect of these non-AI/NLP related technologies outside cloud computing are quite different from the effect of AI-based tools. Individual lawyers have to understand how these technologies operate, but there is no expectation that law firms should guide developments in any specific direction. Lawyers may act as intermediaries in processes that involve the use of such tools as implemented by third parties (like courts, administrative bodies or infrastructure providers), and lawyers have to understand the users' point of view of in relation to such implementation.



### 3. Explaining some basic terminology for AI tools: corpus, datasets, benchmarks and models, linguistic tools and knowledge representation

#### 3.1 Dataset, training set, corpus and training methods

In section 2.1 we have already explained the difference amongst AI, machine learning and rule-based tools. We also mentioned that a machine learning procedure has a *training part* that changes how the machine learning technique actually works, based on a set of examples. This set of examples for training is called the training data, training set or dataset, and the training process results in a trained model. It is the trained model that will be expected to demonstrate the usability of the whole machine learning based tool.

One type of training relies on a pre-recorded set of an input (representing the information based on which the trained model will be expected to make a decision) and a corresponding correct answer, an output value (which should be the result of the machine learning algorithm for that particular input value). This type of training is called supervised learning. Let us say we want a tool to find all the sentences in a contract that are relevant to the duration and term of the contract. In such a case, the training data could be a large body of sentences from real-word contracts, and for each sentence, we have a correct label (sample output) saying whether it is related to the duration of the contract or not. Or a more difficult task would be to categorise court decisions based on certain terms of taxonomy that might overlap (e.g. civil law, matrimonial law, damages etc.). Here, the training set may consist of a thousand court decisions as input and for each court decision, we also have all the expected categories as labels.<sup>19</sup>

Based on the task involved, a sufficiently large training set can provide better results than hand-crafted rules, such as a list of all possible wordings for denoting a term- or duration-related provision in a contract.

Making these supervised training sets is expensive due to human costs, and therefore training sets are often valuable and frequently shared among researchers, and are reused for different training tasks. Sometimes it is not only the full training set with input and expected output that is valuable, but also the input data itself, such as all the published court cases for a given jurisdiction. This latter category of data is called a corpus.

For other tasks, training can also be undertaken without providing any explicit instructions on what is a correct output (no annotation, no labelled answers). These methods are called unsupervised learning methods, and the training process here works differently, because it relies on some previously identified characteristics that make the process suitable for the given task.<sup>20</sup> The designer of the tool will still have to specify how these results should be interpreted for the end-user in the trained model, but there is no manual annotation involved.

One of the most important uses of unsupervised training for lawyers consists of the neural network-

<sup>19</sup> This is called a multi-label classification task.

<sup>20</sup> E.g. based on the methods used, the unsupervised learning process may check during the training how the different text features of the court cases are distributed and provide feedback based on similar court cases.

based language models that have recently come to dominate most fields where NLP tasks take place.<sup>21</sup> Thanks to an enormous amount of information used, such models are capable of capturing subtle relations between the meaning of text in a given language, beyond mere grammatical meaning. But this should not be understood as meaning that the best tools rely only on such unsupervised methods: in many cases, language models such as BERT are used only as a first step, as a pretrained model based on which the designer carries out a further layer of training that could itself be by means of a supervised training method. So, the end-users often access AI-enabled tools which use different trained models building upon each other.

### 3.2 Real-time training and end-user training capabilities

Depending on the actual tool used, AI-based tools might or might not be capable of gathering further data while being used by the end-user.

Some tools are capable of so-called online or real-time training, that is, they may become better and better by being used more and more, but all such tools require some form of feedback from users (to find out if the prediction of the trained model were correct or not). Some type of tools may require end-users to use specific learning functionality (beyond normal prediction work) to improve the usability of the product, e.g. to provide feedback as to whether a highlighted provision is indeed classified correctly as a penalty provision. But many AI-based tools do not have any training or learning functionality available for the end-users at all: they are provided with a specific model as trained by the publisher and the only way to improve the trained model is for the publisher to make an update available to its users, just like the provision of an upgraded software version. Such tools without any end-user training capability are also considered to be AI-based tools.

### 3.3 Benchmarks and claims on the performance of a tool

As we have explained in section 3.1, corpus and training data can be very valuable and is often shared between AI tools.

However, for publishers of AI tools, it is often not enough to share the training data, it is also critical for them to be able to compare how well the various trained models perform, because even if these tools use the same training data, the architecture and parameters are different. The common basis upon which to compare the performance of different tools targeted to solve the same AI/NLP task is a *benchmark*, such as SuperGLUE.<sup>22</sup> Most benchmarks are described in detail in separate, publicly available papers and most (if not all) of such information is an interesting read and understandable not only for IT developers, but also for laypeople such as lawyers. These detailed papers also specify if benchmarks are made up of several, more specific AI/NLP tasks (subtasks).

That is not to say that all AI tools will always have good benchmarks to support the publisher's claim regarding performance of the tool. But there is a way to objectively compare the performance of some AI tools of an analytical nature, and it is to be expected that more and more such benchmarks will come to be used by publishers to compare their respective products.

That is not to say that lawyers are expected to run trained models against benchmarks or to publish results themselves, but we have to understand the use of such benchmarks as marketing tools.

It is important for every lawyer to understand that a promising result for a given benchmark does not mean that a tool using a model built on the same architecture will perform well in real life scenarios. Despite many claims to the contrary, it also does not mean that in a given subtask, computers have achieved or surpassed human equivalence.

<sup>21</sup> For more detail, see Vadász and others (n 6).

<sup>22</sup> See <https://super.gluebenchmark.com> and the [relevant paper published](#) and also the published leaderboard for comparison at <https://super.gluebenchmark.com/leaderboard/>.

*It is important for every lawyer to understand that a promising result for a given benchmark does not mean that a tool using a model built on the same architecture will perform well in real life scenarios*

Also, just because there is an industry-wide benchmark in which a given tool excels, this should not be accepted as a valid claim for the efficacy of the tool as advertised for end-users. At present, industry-wide benchmarks are intended for developers who have practical experience in correctly interpreting the benchmark. If you are faced with a convincing-looking claim for the performance of an AI tool, do check out the details of such claims. Try to read what the advertised benchmark is about before accepting it as true, understand how the given benchmark is measured and whether the use described in the benchmark is sufficiently similar to the intended real-life use. By way of their nature, many popular benchmarks are highly optimised for the given subtasks and datasets, and thus may perform poorly on unseen data from real life scenarios. Most importantly, if you see that, based on the details of the benchmark, an advertising claim for an AI tool is probably misleading, be even more suspicious of other claims made by the publisher.

*If you are faced with a convincing-looking claim for the performance of an AI tool, do check out the details of such claims. Try to read what the advertised benchmark is about before accepting it as true, understand how the given benchmark is measured and if this makes it sufficiently similar to the intended real-life use*

This is not to say that benchmarks are misleading in general, nor that they should not be understood as a general statement as to the usefulness of AI tools in general. Efficacy and efficiency of specific AI tools depend as much on technical factors (like the models and architecture used) as on the specific task at hand, including the users of the tool and the context of usage, economic factors etc. Just because certain tasks are close to each other from a human perspective, will not necessarily mean that a given AI tool will perform similarly in both tasks. This truism applies to industry benchmarks as well, and it is easy to forget this when faced with the ubiquity of misleading marketing communications and the plethora of high-level narratives on the mid-term replaceability of a human work force by AI tools in the legal field.

### 3.4 Why are linguistic (NLP) tools and knowledge representation important for lawyers?

#### 3.4.1. The role of NLP and linguistic tools in AI tools for lawyers

Most of the AI tools that lawyers may use build upon the achievements in the field of computational linguistics, that is, the field of natural language processing. NLP covers diverse areas, including analysing grammatical forms of how words and sentences appear, how these components of our speech and text interact with each other, how they are correctly formed, and how these may appropriately be generated in an automated way.

Research in NLP made it possible to define reliably where the ends of the sentences are in a given text (e.g. does the dot represent the end of the sentence or just an abbreviation?), whether the word “fly” is a noun or a verb in a sentence etc. As long as these simpler tools work reliably a whole new area of automated analysis and generation becomes possible, including finding the subject or the predicate in

a sentence, or the relationship between words in a sentence (e.g. what does the word «that» refer to?).

NLP research relates to the subject-matter described in section 2.1 as well, such as questions as to how the probability rules of mathematics can be used to correct typos in a sentence or to find meaningful relations between a text and a predefined set of categories to describe the text (classification etc.).

If we have to find various proper names (persons, organisations, locations etc.) in a text, we have to apply to the text the appropriate named entity recognition technique defined by decades of NLP research. If we want to find the term of the lease in a plain text contract, we search among the information extraction tools of NLP.

Actually, everything we discuss in this guide will probably fall within the area of NLP research dealing with computational linguistics. Breakthroughs within the NLP field for a language provide a new opportunity that may be of benefit to the use of AI tools for lawyers, and the lack of linguistic tools for a given language will directly hinder the development of AI tools for legal use as well.

### **3.4.2. Legal knowledge representation beyond the surface level of the text**

The type of data lawyers themselves see directly (such as the text of contracts or decisions) is, for computers, unstructured data. That is to say that there is no predefined structure for the data to be processed, it is just e.g. a text file or a video. Prior to further processing and comparison with other data, separate steps are needed to convert this unstructured data into a structure. These steps may include “mining” the meaning of the text, extracting location or terms data of the lease agreement or even just classifying the text as a contract or a decision.

Such a structure can be provided on many different layers (as described above in relation to NLP tools), all at once for the same surface text. These layers may be not only of a grammatical, but also of a semantic, nature, including argumentation (see section 5.3.3) or specific legal meaning.

But we can encode the legal meaning of a sentence (or the legal relations between texts) only if we have a clearly defined way to represent such legal meaning for the computer. This is called legal knowledge representation. Based on the purpose and context of the tools used, it can be as simple as marking documents at the end with a “C” (if the text contains a “contract”) or can be based on very sophisticated legal ontologies themselves building upon many layers of standards, such as a knowledge representation schema for describing air transport passenger consumer complaints based on technical standards.<sup>23</sup>

But legal knowledge representation is far from being a tool for engineers or programmers only. It is of benefit to lawyers to be aware of the existence of such representation languages, and with time, lawyers may be forced to understand and use such a language themselves. For the computer to be able to create this layer of meaning, lawyers will have to annotate plain text for training purposes by the use of annotation tools, constructing the relevant legal meaning in the given knowledge representation language. This knowledge is likely to be relevant for more than only the few selected lawyers who are paid to participate in developing new tools - it can also be of importance to lawyers who want to retrain or fine-tune existing tools to their own specific dataset, as seen in section 3.1.

Knowledge representation can also be used the other way round: to create plain text for humans from a succinct and specialist representation language (as fed into a software program by the user). This is how many document assembly tools work, where lawyers use some knowledge representation syntax to define what text to use in which cases (see 5.1.2). A similar use is evident when lawyers enter all the evidence they want to refer to into user-friendly case management software, where the software captures and enriches relevant metadata of the evidence.<sup>24</sup> Then the lawyers define their lines of

<sup>23</sup> Such as standards W3C Semantic Web standards such as Resource Description Framework (RDF) and Web Ontology Language (OWL).

<sup>24</sup> E.g. by making a reference to an entry on a social media site, the software captures not only the content and the source URL, but also puts a qualified timestamp on the evidence. Or if we refer to e-mail evidence, the software captures not only the content of the e-mail, but also the full message headers, including DMARC results and current name server entries, and timestamps these, enforcing authentication of the headers and the senders etc. If delivery is carried out in registered electronic messaging services, a case management software should record all relevant information retrievable from such registered messaging systems etc.

reasoning in a user interface connecting the pieces of evidence. Instead of having to list and describe each piece of evidence and its characteristics manually (including the sources, provenance, IT details), with this tool the lawyer actually uses the interface to specify what claims each piece of evidence is intended to support and how. This specification by the user when using the interface is actually a form of knowledge representation in the case management software (e.g. visually connecting the pieces of evidence in claims, defining the conditions of applicability etc.). By specifying this information, the user actually gives instructions on how this knowledge representation should be turned into plain text, and therefore how to create a draft of an important part of a court submission (see 5.1.3).



## 4. Opportunities for bars and law societies

Self-regulation of the profession is necessary to guarantee lawyers' professional independence vis-à-vis the state. As AI tools become a more useful part of the tool set of a lawyer, bars and law societies will have to consider if there are appropriate ways to support lawyers in their own jurisdictions.

There can be no single recipe for bars on what to do within the EU, because what bars may and are expected to do depends on different national circumstances e.g. the active involvement of the Conseil National des Barreaux in France in development and research related to AI tools<sup>25</sup> is not necessarily available or even possible for bars in other countries.

For instance, bars and law societies may consider providing useful assistance to lawyers in AI tools by listing or reviewing existing products, contracting third parties to assess or validate claims made by publishers, or even initiating standardisation in such important areas as export and data exchange formats, APIs<sup>26</sup> that could address the serious problems of vendor lock-in and fragmented markets.

Also, bars and law societies will need to continue their active role in defending the core principles of the profession in the interest of clients and the rule of law. For example, issues might arise when AI tools are suspected of interfering with the necessary independence of lawyers (e.g. in the context of platforms which provide referral services)

Small law firms, and sole practitioners who work as contract lawyers for very large clients, might be at a higher risk of pressure to deviate from the core principles of the profession (for instance through attempts at control and surveillance by online platforms), similar to the current challenges faced today by so-called "gig economy workers".<sup>27</sup>

25 Conseil National de Barreaux, Assemblée General du 9 octobre 2020 Groupe de travail Legaltech, '[Legaltechs du domaine de la jurimétrie: Préconisations d'actions Rapport](#)' (9 October 2020) 12 accessed 27 December 2021.

26 Application programming interfaces, i.e. a software interface for connecting different programs with each other according to specific functionalities.

27 See for example challenges faced by contract lawyers in the United States which are in many ways similar to the challenges of gig-workers like Uber drivers: Debra Cassens Weiss, '["Treated like a Robot," Contract Lawyers Chafe under Fickle Facial Recognition Surveillance](#)' (*ABA Journal*, 15 November 2021) accessed 19 December 2021. and '[Managed by Bots: Surveillance of Gig Economy Workers](#)' (*Privacy International*) accessed 19 December 2021.



## 5. Categories of AI tool features and connecting it to lawyers' typical activities

The aim of this chapter is to provide an overall picture of the different ways in which small law firms can use so the many different AI tools at their disposal. In order to do so, these diverse tools will be presented in broad categories that incorporate most of the solutions we think will be used in the future.

We have defined the first three categories based on the functionalities of the tools: drafting support tools, document analysis tools (where the major added value of the tool comes from the documents supplied by the lawyer), and those legal analytics tools where the value is mostly derived from case law or legislation. Chatbot tools and speech-to-text tools are presented in their own categories, and the final category is about tools providing assistance in the internal office administration of law firms.

These categories inevitably overlap with each other, but we have tried to minimise at least the repetitions: e.g. the same technical tools that can be used both for document and case law analytics are explained in more detail only in relation to document analysis, which is discussed first, but that of course does not mean that document analytics is more important per se for lawyers than case law analysis.

### 5.1. Drafting support tools

Products in this category focus on assisting lawyers in drafting different types of documents. This category is divided into two major groups: writing assistance tools and automation of document assembly. A third group is also presented as a generic technical approach for turning data into legal text.

#### 5.1.1. Writing assistance tools

These tools integrate existing text editing and word processing solutions. They augment the drafting process similar to how spell- and grammar-check and autocorrection tools work. These tools rely on the user interface of the aforementioned text editing and word processing tools.

Writing assistance tools can, for example, check the language use from the perspective of what is usually considered desirable language, and word usage as taught in some countries as legal writing or drafting courses (e.g. [Briefcatch](#)). A characteristic of a well-researched tool would be also to provide specific reasons behind the suggestions: what standards are the recommendations based upon, e.g. is it just to ensure less legalese and plain language use (like plain English), or are suggestions based on published linguistic researches on legal texts.

Stylistic recommendations of such tools could also be based on the results of a textual analysis of “winning briefs” or other submissions e.g. if research has shown that there seems to be some correlation

between using hedging<sup>28</sup> text features or positive intensifiers<sup>29</sup> and what kind of briefs were successful. Needless to say, both these types of tools are language specific.

A tool may also assist in verifying whether legal citations are correctly formed and fully refer to the intended cases. When such automated checking of legal citations in the text can be accessed within the editor (as a plugin, e.g. [CiteRight](#) in the US), then these tools are also a subtype of writing assistance. (Often this functionality is available separately as part of document analysis tools, see below).

Some of these writing assistance tools are also geared toward more technical aspects of contract drafting, such as finding undefined terms or empty placeholders and in-text references ([Donna](#)). Another typical solution is to help the lawyer save and reuse text snippets, or even provide some autocompletion based on a personal or firm-wide database (either from contractual provisions or briefs, e.g. [Henchman](#), [xLaw Word](#), [Legau](#)).

Similar to spell and grammar checkers, these tools used to be mostly rule and dictionary based. Recently, and mainly due to advances in neural network-based text understanding, performance of such tools has become better in terms of parsing legal citations (in entity-linking, identifying more citations) and in making stylistic recommendations.

As already mentioned above, these writing assistance tools usually need integration with tools running on the computer of the lawyer (e.g. with Microsoft Word as an add-in) or the tools need to provide the full editing capacity themselves (as an online rich text editor). The advantage of the integration approach is in its ease of use and the short learning curve. These tools will probably remain with us for as long as the main user interface for the work of lawyers remains the word processor. However, for so long as host applications and operating systems change as frequently as they currently do, such an approach makes the maintenance of integration expensive in the long run.

*The advantage of the integration approach is in its ease of use and the short learning curve.*

### 5.1.2. Document assembly tools

The common objective of these tools is to help users automate the assembly of documents from a set of previously recorded provisions and conditions. For lawyers, document assembly usually helps in drafting contracts, briefs and any other documents where automation makes sense due to the repetitive or plannable content.

The designer of a particular document assembly template sets down the business rules that capture the relationship between fixed elements of the text and the external factors to be captured during the finalisation of the document from the template. The authors of templates can be lawyers, but when the emphasis is on more complex business rules, they will normally be external consultants or even professional IT developers working in tandem with lawyers.

These document assembly tools usually work in two phases: (1) the design phase, the creation of the template, also called template authoring, and (2) the use of the templates during a user interview or other customisation or data import process, when the template is turned into specific documents.

Template authoring may be handled on a publicly accessible website (e.g. [Legito](#), [ClauseBase](#)), where the user has to import the text of the provisions from existing plain documents (or type and edit these on the web application), and use the interface on the website to represent the variable parts and conditions of the text that will define the final document, including using business logic to define which provisions are to be repeated based on certain variables (e.g. number of heirs in a will template).

The complexity of the business logic behind the documents is based on the approach of the user. For example, detailed rules can be included in a single lease template agreement which cover lease of

<sup>28</sup> Such as using “however,” “regardless,” or “while” phrases.

<sup>29</sup> Elizabeth Chika Tippet and others, ‘[Does Lawyering Matter? Predicting Judicial Decisions from Legal Briefs, and What That Means for Access to Justice](#)’ [2021] Texas Law Review.

offices as well as storage, and which foresee specific provisions depending on who does the fit-out of the property which is leased etc. This will require defining a more complex business logic. Another approach may be to have different lease template agreements for offices and for storage facilities. This is simpler in terms of business logic, but it makes maintaining the different templates more time consuming, because changes will have to be introduced into two different templates. The optimal choice depends on the capabilities of the tool and the use case of the lawyer. Lawyers are usually interested in reusing as much text (clauses) as possible and defining text as generally as possible. But this comes with a price: the cost of complexity when defining a sample text that can be reused in many different cases.

Web-based assembly tools have to provide the proper interface for defining and integrating the text with all the rules on conditional or repeating text parts (loops, cycles). This may be done purely in a visual, no-code way<sup>30</sup> or by way of domain specific computer code-like scripts.<sup>31</sup> A related function is to connect variables in templates to databases, which enables a user to select, e.g. clients from an existing client database.

This authoring can also be done locally (on-premise), either on the lawyers' computer or on a local server for shared work with colleagues. For these local products, the special authoring functionalities beyond the raw text may be done via, for example, the Microsoft Word application (e.g. [Woodpecker](#) which inserts so-called custom controls in Word documents and a Word panel) or in a dedicated computer program on its own, or even both.<sup>32</sup>

When the template authoring has finished, users may create as many documents from the template as are needed. This second phase is the customisation of a document from the template, which is a) where a user is asked for the individual data for the given document (also called an interview) and/or b) when such individual data is collected from external data sources. Some provisions are defined not based on the direct answers of the users, but as the results of conclusions drawn from the user's choices, which are programmed into the template.

There are big differences in document assembly tools beyond the manner of authoring and operation. This guide treats as of interest only those products which are geared toward lawyers, where authoring is done in a way that lawyers themselves can carry out if necessary. A large number of such solutions are intended to serve a wider, mostly enterprise audience with e.g. supporting the negotiation phase of the drafts or the collection of signatures (e.g. [Juro](#)) or even contract management functionalities such as monitoring expiry of contracts.

This is a relatively mature market, available to most lawyers to some degree since the 1990s, but use of such tools in small law firms is still far from being widespread. One of the major reasons for this is related to the complexities in trying to meet linguistic requirements in more generic provisions and having as few templates as possible. If a law firm uses such a tool to keep a bank of reusable provisions (clause repository), the lawyer expects the template to be able to incorporate and adjust relevant linguistic features (such as declination, conjugation, grammatical gender etc.) of the contract text. This would make it easier for the lawyer to define template texts, but adds an extra layer of difficulty to template authoring.

A very few of the tools provide a certain level of language support for proper verb and noun conjugations outside English (e.g. ClauseBase for French, German and Dutch). While such tool-specific support for smaller languages such as Greek, Czech, Hungarian etc. would be equally important for lawyers working in a number of EU Member States, the problem of market fragmentation already highlighted in the AI4Lawyers project phase 2 makes the provision of such tools and support unlikely.<sup>33</sup>

If the use of such document assembly tools is to gain further market share among lawyers in the future, the abovementioned sources of complexities have to be reduced. Providing easy-to-use, but also powerful, interfaces for laying down applicable rules in document assembly tools is not an easy

---

30 Similar to e.g. [Google's Blockly](#).

31 E.g. HotDocs's scripting language, see [here](#).

32 Such as in HotDocs Author, where Component Studio is a dedicated program that runs in parallel with a Word plugin and the plugin is used for defining simple conditional texts, but the studio is used to define more complex functions and connections to databases.

33 Vadász and others (n 6).

task. Also, providing for easy integration with a wide array of external data sources is challenging. Linguistic problems enhance these difficulties, as the given tools have to provide a precise, tested and reliable output in all supported languages.

If a lawyer user of a document assembly tool is not able to rely on the text which is generated without having to review the results carefully, a lot of the advantage in using document assembly is lost. However, at this level of expected accuracy, it is difficult to create all grammatically correct forms reliably from a generic sentence in plain text drafted by a lawyer.<sup>34</sup> Such accuracy is made possible only if the lawyer defines the model form of the sentences to be reused with special tools (e.g. representation language or with annotations) and not as a normal sentence. Of course, this would mean that lawyers may no longer use plain word processing tools for defining the clauses to be reused. It is possible that, in the future, lawyers will have to rely on more specialised tools even for such simple tasks as drafting new texts, and it maybe that they will require to use legal specific “development environments”.

### 5.1.3. Tools for turning legal data and knowledge bases into text

The final subcategory concerns the generation of text by the lawyer from non-text data into legal texts. Similar to generating readable text automatically from weather data, a tool could create a human readable description of e.g. electronic evidence, a timeline of events, or even create convincing and to-the-point arguments based on the facts and intentions recorded. These tools work similarly to those document analysis tools that create detailed, lengthy reports such as those dealing with certain compliance and due diligence issues. Compared to document assembly, almost no small law firm uses this subcategory of legal drafting support. For small law firms, we consider this subcategory to be only of possible future use, and so in the absence of knowledge as to any widely used tools currently available, we restrict ourselves to painting a picture of the general possibilities in this approach.

The aim here is to turn dominantly non-text legal data into traditional legal writing. With time, lawyers will probably have more and more structured data captured about the cases and documents with which they work, and not just the text surface of the documents with some minor metadata (such as document, case and postal identifiers).

Case management software provides an important source of non-text data and, if lawyers use such solutions more, the amount of data increases and the opportunities to record more data increases even more so. If lawyers record all the evidence supporting their claim or the opposing counsel’s claim, it makes sense to record as much information concerning the evidence as possible, including data about witnesses or metadata such as time-based information (for analysing the flow of events), external identifiers of the evidence and the environment of operation etc. to be used for data mining and further investigations. Evidence based data integrates well also with the arguments and counter arguments used in motions or other court documents.<sup>35</sup>

This means that a sufficiently advanced case management software in a given country provides a good incentive for lawyers to record more and more salient features of their cases contained in the relevant case management system, including the statement of the facts and the related points of arguments to the claim, evidence supporting the claim or the counter arguments to the opposing party’s statements etc. This incentive is further reinforced by the electronic filing guidelines of CEPEJ:<sup>36</sup> mandatory use of smart forms with a number of structured data to be included or requirements on the use of specific citation formats<sup>37</sup> will both lead lawyers to switching from word processors to activity-specific applications.

34 E.g. an expression in Hungarian “the seller shall warrant to the buyer that the electrical safety review of the property has been completed according to Ministerial Decree X» has to be used also for leases, e.g. it should be possible to adjust this into “the lessors shall warrant to the *lessees* that the electrical safety review for the leases has been completed according to Ministerial Decree X». If a law firm wants to retain a unified clause bank as a knowledge base, the document assembly systems they use should be capable of making these kinds of adjustments without the lawyer hand-coding all the possible grammatical rules at the business logic level of conditional texts.

35 See also section 5.3.3 on argument mining.

36 European Commission for the Efficiency of Justice (CEPEJ), ‘[Guidelines on Electronic Court Filing \(e-Filing\) and Digitalisation of Courts](#)’ (9 December 2021) 13 accessed 27 December 2021.

37 E.g., citations of law in European Law Identifier ([ELI](#)) and reference to other cases in European Case Law Identifier ([ECLI](#)).

Having all this data available in the case management system is just half of the story. From the viewpoint of document generation, such a rich database related to a single case can be used to generate complete court documents or important parts of them, with appropriate information in the required format and structured data. As opposed to current document assembly tools, interaction between the human and the tool side is more blurred here.<sup>38</sup>

*Similar to generating text automatically from weather data, a tool could create a human readable description of electronic evidence, or a series of events or even create convincing and to-the-point arguments based on the facts and intentions recorded. ... A rich database related to a single case can be used to generate complete court documents or important parts of them, with appropriate information in the required format and structured data.*

Similar enhancements are possible and already exist outside litigation, such as in contract review or due diligence processes, where lawyers have to work alongside a well-defined playbook and checklists. Document analysis solutions cannot replace humans in checking all points of compliance and need a human review. It may be better to capture compliance information in a review tool used by a human, with the tool then able to create a more comprehensive first draft report or at least certain chapters of such reports. Current language models such as BERT have recently become much better at summarising long texts, including legal texts, and these summaries may also find their way into such first draft reports (see in more detail in the next section, 5.2).

Of course, all such text generation tools would also be language and jurisdiction specific.

## 5.2. Document analysis

### 5.2.1. Introduction

Drawing a firm line between analysis of legal documents (discussed here) and analysis of case law and legislation (discussed in section 5.3) is not always an easy task.

As soon as we turn to analysing submitted documents based on features that have been extracted from case law or legislation, the distinction between analysis of legal documents and that of case law becomes more and more arbitrary.

Document analysis and case law/legislation analysis are both about presenting non-obvious information in legal texts to the users. Further, often the same technology is used for both. If the analysis of the document is based on a machine learning model previously trained to identify 100 standardised contractual clauses, then what is of key importance is the model provided by the supplier, not the document submitted by the user, but this would still be considered as a document-focused analysis tool (e.g. [Kira](#)). On the other hand, if an analytics tool only mines the citations from a brief, and uses these citations as an analysis of how the given document relates to the prevailing case law, this will be a case law focused analysis, and discussed under section 5.3 (e.g. [Doctrine](#), [Juripredis](#)).

However, for didactic clarity, it is better to approach this multitude of AI tools on the basis of two different categories. In this part, we introduce some of the basics of what publishers actually sell us when they say a software “understands” some legal text, and how lawyers can take advantage of such

<sup>38</sup> See section 5.1.3 and also scenario 6.7.

functionalities by using their own documents. This part will try to cover those uses where the lawyer makes the documents available for analysis and the major added value of the tool is recovered from the document itself.

As mentioned above, such a didactic distinction is often ambiguous.

### 5.2.2. Understanding by way of classifying text or parts

The documents to be analysed could be contracts, briefs or other court documents, memoranda, reports – the specific kind of documents to be supplied for analysis will be tool specific, depending on the intended use case and marketing of the tool. For most such tools, the major methods of analysis are related to different approaches to simulating certain aspects of human text understanding.

*The major methods of analysis are related to different approaches to simulating certain aspects of human text understanding.*

Most often the tools classify the document or its parts (at different levels, such as paragraphs, sentences or phrases) according to one or several characteristics. Classification can be aimed at finding the grammatical roles of words in a sentence (is it a verb, a subject etc.?) or at the language of the whole document (English, French etc.) But classification can also be designed to find information on the content of a word (e.g. does a single word or pair of words indicate a private person, a legal entity or a place?), or the content of a document, such as the type of contract according to a statutory or other taxonomy (e.g., Werkvertrag or Dienstvertrag). It is also considered to be classification when the task is to find the existence of a certain clause in a contract (such as a provision on the limitation of liability), or classification can serve to make a suggestion on complex legal issues such as classifying a provision as probably unfair or not.<sup>39</sup>

Based on the complexity and characteristics of the classification sought, this task could rely on predefined rules (such as regular expressions which find matching text) or different types of machine learning-based classifiers trained on datasets.<sup>40</sup> Machine learning based classifiers take into account different features identified in the text, such as document or word meanings<sup>41</sup> or simple frequencies, or even deeper grammatical structures within the text.<sup>42</sup> Classifiers often rely on features identified by the previous runs of simpler classifiers (e.g. first a classification runs to decide the nature of the words in the text, to identify verbs and their related objects, then another classifier is run on these results to see if the meaning of the verb-object pair is related to defining a contractual obligation etc.)

From the developer's perspective, classification tasks are differentiated as being a binary classification task (e.g. is limitation of liability included: yes/no, unfair or not), a multi-class classification task (is the document a contract, is it a memo or is it a brief?) or a multilabel classification<sup>43</sup> (within a contract, label all sentences that contain amendments to another agreement, those that contain data privacy related provisions, those defining the term of an agreement, termination clauses etc., where sentences containing more than one of these labels should be labelled with all relevant labels).

39 Marco Lippi and others, '[CLAUDETTE: An Automated Detector of Potentially Unfair Clauses in Online Terms of Service](#)' (2019) 27 Artificial Intelligence and Law 117 accessed 25 February 2022.

40 Such as those based on probabilities in naïve bayes classification, characteristics identified by training logistic regression classifiers

41 Here, meaning is more precisely meant as embedding, which is a special way to represent documents or words as numerical vectors based on their relationship with other words or documents in the corpus.

42 Part-of-speech analysis or parsing sub-phrases called constituents etc.

43 Don Tuggener and others, '[LEDGAR: A Large-Scale Multi-Label Corpus for Text Classification of Legal Provisions in Contracts](#)', *Proceedings of the 12<sup>th</sup> Language Resources and Evaluation Conference* (European Language Resources Association 2020) accessed 10 October 2021.

### 5.2.3. Analysis based on information extraction: extracting time, relations (citations), content of contractual provisions and facts

Another major type of approach for document analysis is information extraction, where rules are defined or machine learning solutions are trained to extract certain important information from the text, such as monetary amounts, date and points of time in the document or even extracting the text of specific contractual provisions such as a term of a lease contract.<sup>44</sup>

Software tools for information extraction are designed to find and copy relevant provisions from a body of text and use these relevant parts as output, e.g. show it to the user in a report, or use it for further processing, such as comparing the values to a defined threshold and sending a warning to the user if the threshold is not met. Similar to classification, the task of information extraction can be based both on rules (including regular expressions) and on machine-learning based methods.

This extracted information may be as simple as just an amount of money or the name of persons or legal entities mentioned in the text. Based on sufficient dataset and training, information extraction can also be used to extract accurately sophisticated relationships in the text, such as extracting from witness statements persons and their birth dates or places, or a person and his job title, a legal entity and the employment relations, affected locations etc. (Needless to say, this use case is mostly relevant when a large amount of documentation has to be processed: it is not going to find new information that a diligent human could not find.) Information extraction works well also with time related information (obtaining dates from a text in a standardised format, turning “yesterday» or durations such as “four days» into computer processable format) or identifying typical events in long text, even outlining the temporal order of the events so identified. This NLP task may also be used to find out to what or whom certain pronouns relate (called coreference).

When a tool goes through the document to extract all the citations (for example to check whether they are correctly formed and point to the intended cases) that is also information extraction (e.g. WestLaw’s [Quick Check](#) for US cases.) Such extracted data can be loaded into case management or contract management software or reused for due diligence or compliance purposes.

### 5.2.4. Combining classification and extraction for document understanding and analysis purposes

Classification and extraction tools often rely on each other, as an extraction tool is trained and run only on those provisions that have been classified as relevant (e.g. a classifier decides if a contract is a lease contract, another one labels those provisions in the lease contract that contain lease amounts, and an extractor extracts and presents the amount to the user). When whole provisions are extracted as relevant, a different type of algorithm may compare such extracted provisions to a “desirable” (e.g. legally approved) provision and score such differences based on similarity in meaning (embedding) or just present the compared version to the user to decide on the degree of risk.

These tools are always language specific solutions and, if the analysis relies on legal concepts, that will make it jurisdictional specific as well. Although some products are marketed as language agnostic (e.g. in document classification), one can expect that using it for languages other than English will negatively affect the performance of the tool.

For lawyers, usability of such a tool depends heavily on the specific use case which the given tool was intended to address, and if a lawyer wishes to work with tools for non-English languages, there will be very few such tools offered.<sup>45</sup> Some tools are intended to assist in the more thorough overview of a large body of documents based on a predefined set of questions, whereas others are optimised for reviewing one specific document at a time.

Most of the tools work with rules and machine learning models that the supplier defines in advance, and only the supplier is able to change these, but there are also tools that make it possible for the user to train the tool to identify new provisions (such as identifying new risky clauses, red flagging them in due diligence software, e.g. [Kira](#), [Luminance](#)).

<sup>44</sup> Ilias Chalkidis and others, ‘[Neural Contract Element Extraction Revisited](#)’ (2021) abs/2101.04355 CoRR.

<sup>45</sup> Conseil National de Barreaux, Assemblée générale du 9 octobre 2020 Groupe de travail Legaltech (n 25) 42.

The result of combining classification and extraction could be to produce, for example, a report on the existence of certain data in the document, or an evaluation of compliance or certain predefined categories of risks, or even a probability based “prediction” of consequences (like the need to mitigate a certain risk). Some tools even provide a commented Word document with suggested tracked changes (provided the use case is reviewing an English or Dutch language data processing agreement or a non-disclosure agreement, e.g. [Lynn](#) relying on a clause identification engine and on domain specific legal knowledge embedded in the tool by the supplier).

*... the majority of such commercial tools are cloud-based tools where lawyers have to upload documents to be analysed, which makes frequent updating of the model by the supplier easier. However, for lawyers as users, this naturally raises the question of confidentiality and access to the results and history of analysis.*

The deployment mode of the tools and user interface are tool specific, but the majority of commercial tools are cloud-based tools where lawyers have to upload documents to be analysed, which makes frequent updating of the model by the supplier easier. However, for lawyers as users, this naturally raises the question of confidentiality and access to the results and history of analysis.<sup>46</sup> Considering that the usual result of a document analysis is a text or report, the problem of exportability is less relevant for this category of tools (compared to e.g. document assembly tools).

### 5.3. Text retrieval and analysis of case law and legislation

#### 5.3.1. Introduction

Computer assisted legal research or text retrieval has been with us since the end of the 1960s<sup>47</sup> and has been in widespread use in EU countries at least since the early 90s. Traditional techniques of indexing legal texts, and the use of professional citation systems for case law, have also been adapted to the computer environment and have thrived ever since. Optical character recognition and constant broadband access make it commonplace for us to access even historical legal records.

Nevertheless, these reliable techniques of indexing and search have proved to be insufficient for many purposes, especially given the development of new requirements. We want to access more and more information, but in a shorter time: a full-text legal search based on keywords may be sufficient to find a specific term in a single book, but it is inadequate to find a more common term in appropriate context among one hundred thousand judicial decisions. Having too many results is not much better than having no result at all, so the focus of text retrieval changes to obtaining more relevant results.

Relevance and the expectation of finding out the “appropriate context” of the search also leads to other major changes (see in section 5.3.3 in more detail). First, we need tools that can accurately translate the queries as formulated by lawyers into something that enables precise database searches, in such a way that lawyers will not need to be trained for weeks just to use the tool. Also, different tools are needed to find hidden structure and information in legislation, case law and other legal texts. This hidden information is uncovered by analysing these legal texts. If lawyers want to be able to formulate queries like “show me all the cases where a business secret claim was upheld despite the secret being reverse-engineerable”, this assumes that a sufficient number of such cases has been properly analysed beforehand with such questions in mind.

<sup>46</sup> See section 7.3.3.

<sup>47</sup> Charles P Bourne and Trudi Bellardo Hahn, ‘[A History of Online Information Services, 1963-1976](#)’ (Cambridge, Mass : MIT Press 2003) 244 accessed 27 December 2021.

New developments in machine learning and natural language processing have made it possible to provide solutions to these changing needs. These same techniques make it also possible for more lawyers to access a wider range of analytics and jurimetrics of legal texts (especially in case law), which was simply not feasible previously or which was very much limited in scope due to the high costs of manual labour needed. Costly manual annotations were previously a prerequisite for utilising many machine learning solutions, but NLP techniques (especially the so-called foundation models<sup>48</sup>) have made it possible in many fields to substitute or at least minimise the need for human annotation and so open up new ways of using machine learning.

As with document analysis, it is impossible to show all the possible ways such legal texts can be analysed, but we discuss three distinct matters: first, the differences between advanced text retrieval and search engines with quantitative data and what are often called predictive tools, second, provide an insight into the futuristic area of how argumentation mining could be used by lawyers, and finally we briefly mention tools which provide quantitative data on the activities of certain participants in court and similar procedures based on case law.

All such tools of legal analytics depend on one important prerequisite: a sufficient volume of legal texts in electronic format available for analysis, including legislation or case law. It is also important to highlight that some entities may have access to legal information that is vastly superior in terms of size or quality compared to that to which other entities have access. In the field of legal analytics, it really matters whether the analysis of a case is able to rely only on selected and manually anonymised court decisions in published and edited form, or on the full court file similar to the “Public Access to Court Electronic Records” (PACER) in the United States. This is especially important in terms of how reliable predictive tools can become: training tools predicting the outcome based on the judicial statement of facts or summaries of claims as included in the decision itself can be less effective than those based on the full texts of filed claims.<sup>49</sup>

Those having access to more information will be at a considerable or even critical advantage compared to other entities. Such privileged entities could be courts or other government bodies (such as the prosecution service, tax authority or police), but when compared to small law firms – which can only access publicly available services based on publicly available information (and by nature have access to a much smaller corpus of information) – even larger commercial entities (such as insurance companies or large law firms) may have a considerable competitive advantage.

*Those having access to more information will be at a considerable or even critical advantage compared to other entities. Such privileged entities could be courts or other government bodies. Moreover, when compared to small law firms which can only access publicly available services based on publicly available information and have by nature a smaller corpus of own information, even larger commercial entities (such as insurance companies or large law firms) may have a considerable competitive advantage.*

(For the distinction between legal analytics and document analysis, see section 5.2.1.)

<sup>48</sup> See footnote 113.

<sup>49</sup> Masha Medvedeva and others, ‘Automatic Judgement Forecasting for Pending Applications of the European Court of Human Rights’ (2021) 4 accessed 25 February 2022.

### 5.3.2. On three different objectives of legal analytics

Thanks to the new technologies of NLP and machine learning, the offerings of legal text retrieval services have vastly diversified. It is no longer just the raw text of the judicial decision or the legislation that is displayed to the user. More and more information is searchable besides the keywords in the text. This includes such information as the court making the decision, legal areas, results of the decision or the date of decision. Relations between decisions are presented, so that one can follow the path of a judgement through appeals, see how many other judgements have cited this exact decision and whether they yielded the same results or not. A report of the Conseil National des Barreaux calls this the level of “informative justice”.<sup>50</sup>

But it is not only the name of the court and the date of the decision that can be extracted from a legal text. The next level of exciting new features is based on the extraction of any other quantifiable data, like the sum of compensation awarded in personal injury cases, child or spousal support payments granted in family law, amount of fines imposed by a competition authority, or even compensation figures in common disputes concerning residential leases between the landlord and the tenant, or labour law indemnities to be paid in case of unfair dismissal of an employee etc.

When having such quantifiable data on all these decisions, it often makes sense to search and record further related information. These frequently mined features provide essential context for the decisions: the type of damages suffered for personal injury cases, the income figures of employees affected by the termination, or of their spouses etc. Of course, whether any such data can be recovered from case law or not depends not only on shrewd machine learning algorithms, but on the data available in the published cases. Considering that a number of common law court decisions tend to be longer with more citations and references than decisions in similar cases in continental law countries, this also has an important effect on what can be achieved in legal analytics. Requirements of anonymisation and interpretation of the same regulation on data protection are still very much country specific and these also affect the opportunities of NLP and machine learning. Even within the same country, information included in published cases depends on the area of law affected, and on regional or local customs, or the individual styles of judges.

Through the use of such analytical tools, case law becomes more transparent in terms of quantifiable information, and this is what the Conseil National des Barreaux calls “analytical justice”. Analytical justice focuses on making past cases visible to users through queries based on figures.

There is also a third level, where – having been trained on historical data – an AI tool makes estimates, forecasts or predictions in terms of quantifiable information for a given case. How many years of prison sentence will a person probably receive based on such and such a statement of facts? Or what is the probable amount of damages that a person injured in this way will be awarded, and so should they accept the offer of the insurance company for € X?

All such AI tools for forecasting or estimation are likely to be a valuable addition to a lawyer’s practice. In English speaking countries, these tools are often called “predictive justice” tools. Some consider this term to be misleading,<sup>51</sup> but most of the suggested replacement terms are not really much better,<sup>52</sup> and so we do not attempt to settle this dispute here. One has to be aware that prediction has a specific meaning in machine learning parlance, simply meaning the output of the model trained on historical data. In that sense, even ranking of legal texts retrieved from a database in terms of relevancy is a prediction in itself (prediction of relevancy), in contrast to which predictive justice in this context usually means an output on the expected terms of the judgement or the outcome of a court process based on historical data.

We do not label these three levels as levels just because level “two” or “three” (analytical or predictive justice) would be more advanced in many ways than level “one” (informative justice). The label does not even mean that the levels become more useful for lawyers or more disruptive. We call them levels

50 Conseil National de Barreaux, Assemblée générale du 9 octobre 2020 Groupe de travail Legaltech (n 25) 63.

51 *ibid* 9.

52 We avoid using the term „simulative justice” as the third level, as suggested in the report from the Conseil National des Barreaux already mentioned (*ibid* 63.), because that term is based on how the working of a specific tool was explained to drafters of that report, and it would be misleading to use the same term for other tools. Other suggested terminology includes e.g. „forecasting” of decisions.

because it highlights the historical process of how solutions in informative justice are a prerequisite for analytical justice, and access to quantifiable figures in case law is in turn a prerequisite for any machine learning based divination or dispute resolution algorithms of a predictive nature.

However, one has to be mindful that a predictive justice AI tool comes with its own considerable risks and special dangers compared to the other two levels.

In informative and analytical justice, interpretation is undertaken by the user (the lawyer), and there is no risk of introducing further bias into the service provided (other than the bias already included in past cases, but that is rarely the responsibility of the lawyer). Today, lawyers are not expected to understand how text retrieval actually works and why certain results might have been excluded. If that missed result negatively affects the service provided by the lawyer and the outcome of the case, the lawyer will nonetheless bear full responsibility to the client according to current rules of deontology.

But at the third level of legal analysis, due to its new functionalities, the lack of explainability and introduction of new bias could cause a problem, and therefore lawyers using such tools should be aware of the increase in such risks and take appropriate steps if necessary. Both the ethics guidelines of the European Commission<sup>53</sup> and the draft AI Act<sup>54</sup> define transparency requirements for such tools (even though AI tools are not currently considered as high risk so long as they are used by lawyers and not by judicial authorities<sup>55</sup>).

*A predictive justice AI tool comes with its own considerable risks and special dangers compared to the other two levels [informative and analytic justice]. ... The lack of explainability and introduction of new bias could cause a problem, and therefore lawyers using such tools should be aware of the increase in such risks and take appropriate steps.*

It does not mean that forecasting tools should never be used unless all the business secrets behind the internal workings of the tool are revealed. However, for as long as a lawyer has not even a general idea as to why and how the tool suggests certain results, this lack of understanding and explanation should be clearly disclosed to clients and should be taken into account by the lawyer during the particular mandate for the client.

Even today, one can find many automated dispute resolution mechanisms which work without any explanations about their automated decisions. If clients are aware of the use of such predictive black boxes, e.g. as a mutually acceptable starting point for a negotiation or mediation with a counterparty on the amount of compensation, then this could still be something worthwhile considering.

It is important for lawyers to try to evaluate these tools according to their weight and usefulness and this requires some level of information on how reliable are the suggestions made by the tools. To understand the working of such tools to the necessary degree, lawyers will probably need both special training and also the ability to access specialist resources if needed. Transparency and explainability do not mean that explanations should be accessible and understandable by everyone, including non-technical people. Should explanations go too far in avoiding technical terms, this would result in explanations which are not able to provide transparency for anyone, not even for the purposes of audit and review by specialists. Further, broad explanations that build too much on analogies between AI and actual persons or real-life processes can be seen as dangerous, due to transferring trust in certain roles and processes in society to AI tools.<sup>56</sup>

53 European Commission, '[Ethics Guidelines for Trustworthy AI](#)' 18 accessed 12 December 2021.

54 '[Proposal for a Regulation of the European Parliament and of the Council Lawing Down Harmonised Rules on Artificial Intelligence \(Artificial Intelligence Act\) and Amending Certain Union Legislative Acts](#)' accessed 12 December 2021.

55 See Ibid Article 13 and Annex III section 8.

56 Like explanations on how "virtual judges" make decisions in a simulation where virtual judges are actually just decision trees trained on dozens of features in hundreds of cases.

*To understand the working of such tools to the necessary degree, lawyers will probably need both special training and also the ability to access specialist resources if needed*

### 5.3.3. Advanced searching techniques beyond the text: semantic search and argument mining

Even what is called “informative justice” in the previous section has a great deal of potential for improvement, and advances in this area may fundamentally change how lawyers work in the future.

Semantic search is a new technique. It is not a specific technique, but a multitude of techniques with similar objectives. The objective is to find the best way to retrieve text based on the intended meaning of the query, i.e. a search that relies on layers of meaning beyond the surface of the text. With this approach, a search solution which can provide the most relevant results by finding the appropriate context of the text will perform better. It should find synonyms of a given term (which is of course a domain specific issue in law as the same words may have very different meanings in other areas of law), but should exclude results that are not related to the intended purpose of that given search (for example, if you are looking for civil law cases of defective performance in the case of second-hand property, then it should not return cases dealing with tax-related issues of second-hand property.)

These semantic searches are nowadays typically implemented in two ways: one is when the user provides a sample text and the solution is to search for similar texts in the database. The other approach is to ask a specific question and have the tool find some cases and answers in the database relevant to the question (question-answering).

In this sense, argumentation mining and searches based on argumentation are one form of implementation of the objective of semantic search. It is based on NLP tools finding argument structures in case law, such as what kind of units of arguments exist in a given legal text, which expressions are conclusions, which are premises, which expression is just a logical connective, and whether one unit of argument supports or refutes another etc.

Argumentation mining is a very well-established multidisciplinary field in AI research, having originally had the intent of computerising the even older field of rhetoric, aiming to help find the correct arguments with the help of computers. This involves finding arguments in natural language texts, finding typical schemes of argumentation and visualising argumentation, but also finding arguments in a given field and assisting in the verification of lines of arguments or making summaries of argumentation. The field of law has always been a good domain for carrying out such AI research, due to the large number of available data with sufficiently high-quality arguments.<sup>57</sup>

This area is very promising for lawyers as well, because the tools can help us find lines of reasoning present in briefs and judicial decisions written in natural language. It can help uncover deeper structures in judicial decisions, such as what kind of evidence is needed to prove a statement of a particular fact. Even in continental systems with complex statutory provisions and codes in place, this is rarely a trivial question, due to the difference between law in books and law in action, and also considering that many detailed particularities cannot be specified in legislation.

<sup>57</sup> Prakash Poudyal and others, ‘[ECHR: Legal Corpus for Argument Mining](#)’, *ARGMINING* (2020) accessed 25 February 2022 and Marco Lippi and Paolo Torroni, ‘[Argumentation Mining: State of the Art and Emerging Trends](#)’ (2016) 16 ACM Transactions on Internet Technology 1 accessed 25 February 2022.

*The [argumentation mining] tools can help us find lines of reasoning present in briefs and judicial decisions written in natural language. It can help uncover deeper structure in judicial decisions, such as what kind of evidence is needed to prove a statement of a particular fact*

Based on large amounts of manual work, a number of interesting systems have already been built that can demonstrate the power of these tools. For example, in a system specifically made in the US to cover business secrets, researchers identified, in the case law 27 different factors treated as relevant in decisions made by judges, as being factors which weakened or strengthened the arguments and the position of the claimant etc.<sup>58</sup>

Of course, these systems were based on expensive manual work, but in theory, with such factors identified along with their effects on the judgement, and with other structures of arguments uncovered, AI tools could provide further assistance to lawyers. What are the weaknesses of my planned claim? Where is there a weakness in the chain of reasoning of the opposing party's brief? Based on case law, what do judges accept as proof of an unfair dismissal of an employee on sick leave? How can we justify that this landlord acted as a fair and reasonable one, etc.? If we are able to find such arguments in case law directly, this would be a more effective search than trying to find the same arguments based on keywords only.

If we are able to create a reliable representation of argumentation in case law, that could enhance not only informative justice, but also other levels of legal analytics.<sup>59</sup>

#### 5.3.4. Analysis of activities of participants based on case law

The last subcategory of AI tools analysing case law is the one which deals with analysing the records of different participants in a court procedure, usually judges and lawyers, but in some countries also the claimant and defendant.

In the United States, with the availability of PACER and a similar wide range of court documents, detailed information can be found from career information on judges to motion types decided and case histories, with a very fine level of breakdown in many categories, from grant rate to duration and how these numbers compare to average figures. In the US, analytics on lawyers include clients the lawyers have represented, their court appearances and the types of cases they have worked on, and of course, cases won and lost and trends etc. In France, where analytics on the individual activity of judges is prohibited by law,<sup>60</sup> lawyer analytics still provides basic biographical information on lawyers and the number of procedures in which they have participated, including the subject matter of the procedures in which they have experience (their expertise) and even the name of their business entity clients.<sup>61</sup> But in Hungary, where every name is anonymised in court decisions other than the names of judges and lawyers, obviously the identities even of business entity clients will not be available for analysis.

As one can see, the exact scope of analytics depends on the information available on such persons in the given country and the degree of anonymisation generally applied in case law. These factors therefore play a large role in determining the usability of such AI tools. Apart from technical questions, the analysis is connected to a number of serious ethical issues and, for example, can impact on the independence of judges and the rule of law.<sup>62</sup>

58 Kevin D Ashley, *'Artificial Intelligence and Legal Analytics'* (First, Cambridge University Press 2017).

59 Case Law Analytics in France already provides an interesting estimation: in a given hypothetical case defined by the user along a number of parameters, it provides some estimates as to which factor has what effect on the chances of winning or being awarded certain sums.

60 *'Article 33 - LOI N° 2019-222 Du 23 Mars 2019 de Programmation 2018-2022 et de Réforme Pour La Justice (1) - Légifrance'* accessed 28 December 2021.

61 Subject to the rules on anonymisation of decisions.

62 Council of Bars and Law Societies of Europe, 'CCBE Considerations on the Legal Aspects of AI' (n 3) 17.

## 5.4. Speech-to-text tools

Converting live speech and verbal instructions and commands into written text is important for legal uses as well. Reliable speech-to-text software for such conversions already exists for probably all official languages of the EU.<sup>63</sup>

That means probably the most important use of this software for lawyers, transcribing dictation to text, is already possible for these languages. There are still important differences amongst the tools and how suitable they are for professional use: the recognition error rates may be quite different, and the ability to train the tool on the specific speech patterns of an individual, for specific vocabulary, can also be of practical significance. Otherwise, this specific use of AI is already more a question of the commercial aspects than of technical availability.

Of course, besides dictation, there is still a number of uses with ample room for technical improvement. This includes transcribing noisy speech or unspecified persons (e.g. for depositions), or identifying multiple speakers and converting such real-life conversations into text despite the cross-talk.

The task of recognising the speech of unspecified persons is made considerably easier if the range of distinct input to which the tool is expected to react is limited, such as when using voice assistants that are expected to distinguish between a couple of dozen commands. In theory, this makes voice recognition tools a practical addition to a law practice when combined with chatbots, discussed in the next section.

With the simple transcription of spoken text to written text, much information present in the spoken voice is lost, such as emphasis added by stressing certain words, or differences in tone or rhythm. This is often critical for understanding the spoken text and, without it, a simple transcription can be almost useless. Therefore, extraction of such meta-textual features from speech is also of great importance for transcription tools. Having such tools is a prerequisite for further innovation, such as automatically creating the minutes of the general meetings of a company from the audio recordings thereof.

Whether for dictation, transcription or for voice assistant purposes, before implementing these AI tools, lawyers should become acquainted with the way the voice of the caller is processed. This data is personal data and is probably also client data that is subject to professional secrecy obligations, requiring specific consent from clients.<sup>64</sup> Currently, most speech to text tools, even those dedicated to dictation, work only via an internet connection, meaning that they transfer at least some pre-processed form of the original voice (which is still personal data and is subject to professional secrecy). Until a number of years ago, the mainstream deployment model for dictation tools used to be on-premise, but now such on-premise use is often a premium product or reserved only for larger businesses.

[box: Before introducing these AI tools, lawyers should become acquainted with the way the voice of the caller is processed. This data is personal data and is probably also client data that is subject to professional secrecy obligations]

<sup>63</sup> Homoki (n 5) 38.

<sup>64</sup> Prior to using voice assistant services for client work, lawyers should familiarise themselves with the interpretation of the European Data Protection Board in their guideline, as lawyers would probably also be considered as data controllers, see European Data Protection Board, '[Guidelines 02/2021 on Virtual Voice Assistants](#)' (7 July 2021) 16 accessed 14 January 2022.

## 5.5. Chatbots

Chatbots are dialogue or conversation systems that simulate human conversational capabilities to a certain degree. The exact functionalities that chatbots provide change with time and user expectations. Since 2008, the growth in the use of social media and messaging applications has been very strong, and now 57% of the EU's population uses these channels.<sup>65</sup> No wonder that with such growth in these channels, the popularity of using chatbots has also grown, because chatbots are excellent tools for providing an interface between businesses and users on these platforms. With appropriate integration, the same chatbot can serve many messaging applications and social network sites – although differences in focus between apps/sites often necessitates divergent chatbot features.

Chatbots may appear integrated into websites of the firm or at a contact point of the lawyer on social media or messaging platforms (e.g. on a Facebook page or a messaging bot etc.). Although there are already chatbot platforms that offer no-code solutions, integration and implementation will usually require the involvement of specialists.

When using chatbots, interactions by the users are mostly managed by the user either choosing from a set of possible menu options or typing in a question. (In the case of integration with a speech-to-text solution, this could also be managed by speaking to a virtual assistant, but this is currently not a frequent use case for law firms.) When a user can choose only from a fixed set of buttons, the experience will be quite similar to that of a website with menus. Of course, the more difficult approach is the one where the chatbot has to guess the intent of the user from the typed-in message. Usually, this will require natural language processing capabilities that are specific to the language of the conversation (although simple patterns are often used in some simple cases). There are powerful chatbot platforms capable of “understanding” all EU official languages and more, but some of the easiest-to-use and most-popular international chatbot platforms are still limited to natural language understanding in English. To circumvent that limitation, chatbots are sometimes integrated with translation tools (DeepL, Google).

In the case of typed messages, natural language understanding here means that the chatbot platform tries to guess the general *intent* of the user (what the user wants the software to perform) and some relevant highlights from the message, such as any names, location, time or date etc. entered (“*entities*”). This guess is based not only on the message itself, but on the training previously implemented by the creator of the chatbot, and also on the prior conversation history with that given user. The intents and entities to be identified in the message will have to be predefined by the creator, but the actual method and effort needed for such definitions depends on the platform used. Based on the intent and entities discovered, the chatbot platform will either provide a predefined response to the user or turn to some external integration (such as to book an appointment with the lawyer).

While the opportunity seems obvious, many typical use cases of chatbots are not a good fit for the average small law firm. Law firms of this size are rarely able to provide direct sales opportunities for customers, such as through their web shops, and it is difficult to generate valuable leads through the simple interactions of chatbots. Also, tight integration of a chatbot with live human operators is not something that a typical small law firm can afford.

*While the opportunity seems obvious, many typical use cases of chatbots are not a good fit for the average small law firm.*

Yet, there are a couple of secondary areas where chatbots could be a useful addition for small law firms. Importantly, one should keep in mind that any of the use cases set out below could be in breach of national deontology rules, which is not a criterion that this guide can take into account.

Theoretically, chatbots could be used to provide static, mostly predefined legal advice to clients, or even start an external document assembly system to create a complex document for the user based on the

<sup>65</sup> See [EUROSTAT table ISOC\\_CI\\_AC\\_I](#), accessed 29 December 2021.

“interview” carried out in the chatbot (see section 5.1.2). However, that will probably never be a very popular and practical approach due to the risks involved, even if the legal advice is based only on the user choosing from menu options. Based on the current reliability of natural language understanding tools, trying to guess the legal needs of a client based on the client’s typed input is not yet a good idea, even in English, let alone when using two-way automated translations for understanding and response.

Chatbots can be integrated into reliable online identification tools, which could make it possible to use chatbots to serve existing (probably individual) customers and provide them with information regarding their account, supplying information about unpaid items or the status of their proceedings or any service being provided by the lawyer. Of course, currently this functionality is not common at all, reliable online identification tools are not readily integrated with chatbot platforms, which means that customer service by chatbots will be limited to providing the contact details and opening hours of the law firm, and also possibly booking an appointment.

Regardless of any difficulties in sales, chatbots can also be used for marketing purposes by law firms. Besides presenting chatbot users with the website, the blog and any other online presence of the law firm, the organic opportunity here is to make it possible for users to subscribe to customer lists of the given platform or as followers, depending on the platform (e.g. for the old-fashioned ones, to subscribe to e-mail newsletters).

Obviously, creating a customer list from a chatbot interaction is subject not only to deontology rules, but also to considerable data privacy requirements requiring consent and privacy notices.<sup>66</sup> If the law firm is intent on gathering further details of chatbot users as recorded during discussions or on saving the content of the chat as well, consent and privacy notices should clearly cover this information.

As a matter of privacy, should a law firm decide to use a chatbot to provide a client service or to otherwise include client information during the chat (whether from a question or in a response), professional secrecy obligations will have to be complied with, including transfer of personal data to third-party processors, probably also out of the European Economic Area. This is a particular difficulty, as clients can provide just any information to the chatbot, including information that would be covered by professional secrecy – even if this was never the intent of the deploying law firm.

In summary, chatbots can serve as a foundation for many online channels to keep in touch with customers, or even to ensure a presence in a virtual world.<sup>67</sup>

Chatbots cannot be used in a vacuum: the stronger the online presence of the firm, and the more the firm spends on online marketing, the more it makes sense for that firm to invest in chatbot tools as well. There is not much point in creating a chatbot for the sole purpose of booking an appointment with the law firm.

The untapped potential of chatbots seems considerable, but the barriers to their adoption are not technical in nature.

---

66 See Fashion ID judgement of the European Court of Justice (ECLI EU:C:2019:629) on a platform provider and the group administrator both considered as a joint controllers.

67 In late 2017, the Second Life Bar Association closed its doors. That only shows that similar initiatives we now see in Decentraland already have significant precedents.

## 5.6. Assistance in internal office administration by AI tools

This section is an umbrella-section where we introduce some features of AI tools that do not fit well into the other sections above, but could still be helpful in increasing the internal efficiency of law firms.

One report concerning small and mid-size US law firms shows that an average of 31% of the working hours of a lawyer is billable work (utilisation).<sup>68</sup> It is also worthwhile noting that, based on this report, this figure does not seem to grow with time, despite increases in law firm automation. Earlier reports show that a considerable ratio of such non-billable time (48%) is spent on administrative tasks such as office administration, billing, configuring technology and debt collection.<sup>69</sup> Obviously, with the help of automation, the time spent on some of these activities could be further reduced.<sup>70</sup>

As discussed in more detail during phase 1, most such automation does not necessarily need to involve AI tools at all.<sup>71</sup> Classic business software for law firms, usually called practice management tools (or ERP, CRM etc.), should be the first step in decreasing the non-billable time necessary for running a law firm. A major dilemma is that if lawyers even now can spend hours on troubleshooting technology and transferring information between software, organising electronic workspaces, etc.,<sup>72</sup> then using more and more software that does not integrate well might even increase the number of human hours needed to fix IT gaps. The issues of reliable integration amongst software programs are more questions of appropriate software architecture and implementation, developments and consultation, than of using a new tool with some claimed artificial intelligence capabilities.

In this section, we limit the discussion to those tools in this field that rely on artificial intelligence functionalities.

The first major area where AI tools can be of help is, of course, surveillance of lawyers – less controversially called time tracking. Efficient software exists which can automatically record all time spent on a desktop, a laptop or a mobile device, based on the name of the applications used and the documents the lawyer worked with, which can rely on metadata saved in the document management system or the e-mails sent while undertaking the work. It is also easy to track and extract location details from a mobile device for court proceeding purposes and so match calendar entries with timesheet entries. Based on this information, the tool assists lawyers in categorising their activities according to matters, or even to some more refined item on the invoice. If the billable activity is tracked properly, that also means that any documents saved during this activity can also be saved and filed to the same matter (folder etc.), without the lawyer having to enter client data both for time tracking and document management purposes.

Before introducing these tools, a law firm should clearly define what it wants to achieve with this tracking software, and implement only software that does not go much beyond the defined limits. Because these tools are language agnostic and not jurisdiction-specific, one might easily be tempted to buy the latest and most popular product, with attractive introductory prices. The risk is that many employee-monitoring software programs are advertised under the much more acceptable name of employee productivity- or time-tracking software. However, the capabilities of the software go well beyond the expectation of a lawyer, including total worktime surveillance with the possibility to capture continuously the content on computer screens.

68 Themis Solutions Inc., '[2021 Legal Trends Report Published by Clio](#)' (August 2021) 34 accessed 30 December 2021.

69 Themis Solutions Inc., '[Legal Trends Report 2017 Powered By Clio](#)' (2017) 13 accessed 30 December 2021.

70 For a similar comparison of what was the distribution of billable hours in a survey in the US, see Homoki (n 5) 31.

71 *ibid* 42.

72 Themis Solutions Inc., '[Legal Trends Report 2018 Powered By Clio](#)' (2018) 14 accessed 30 December 2021.

*The risk is that many employee-monitoring software programs are advertised under the much more acceptable name of employee productivity- or time-tracking software. However, the capabilities of the software go well beyond the expectation of a lawyer, including total worktime surveillance.*

Many time-tracking software programs also provide reports on individual users in a dashboard style, and with each new version, they try to become more attractive by expanding their functionalities with evaluation of the user (identifying outlier employees in terms of billable hours) or security features, such as data leak prevention measures (recording information on exported documents and emails etc.).

Even those time-tracking software programs that do not go beyond actual time-tracking functionalities are now more often than not cloud-based, involving sub-processors, possibly outside the European Economic Area.

So, before introducing automatic time tracking software, professional secrecy and privacy implications should be clearly evaluated, including a legitimate interest assessment, checking any data exports outside the EEA. Moreover ethical and labour law concerns need to be addressed, including ensuring proper notification of users.

Besides time-tracking, surveillance is becoming a daily part of lawyers' lives, with every mobile and desktop operating system providing biometric authentication of users (including facial recognition, of course, without the possibility for the law firm to access biometric template data). Also, not only security cameras, but video doorbells have facial recognition capabilities which could theoretically be deployed at law firms as well. This guide does not take sides as to whether law firms should use such technology for security purposes or not, but it is important to highlight that carrying out surveillance by using AI is not only an issue for governments, it is within the technical capabilities of private persons and the smallest of businesses.

AI tools also have other uses in law firms beyond surveillance. One of the most useful are those analytical tools that help the review, standardisation and streamlining of the time recorded by the lawyer: they highlight double entries and probably similar entries with different descriptions, missing narratives or other inconsistencies within the logged information. Even individual client expectations can be configured, where alerts will be raised in case of non-compliance with such expectations. And lawyers should not be surprised that, even if they do not use these review functionalities themselves, a client might still use them, because the same tools also work for incoming invoice data and narratives.

Besides the analytics of time recorded, AI tools can also be of great help in document management for finding and suggesting the appropriate matter and metadata of incoming and outgoing communications based on communications and work history of the user or other predefined patterns.

Such tools also exist of course in areas of accounting and book-keeping. Even simple tools may help in matching cost items in bank statements with supplier invoices or with matter management (for invoicing it to customers as expenses), or label documents according to specific categories that help streamline the book-keeping work. Naturally, the more information that is recorded with regard to the financial details of a law firm, the more possibility there will be to monitor and analyse the financial situation, including controlling costs and profitability - but these tools go beyond the subject matter of this guide, which is intended for small law firms.



## 6. Scenarios

### 6.1. Introduction

In this section, we give some easy-to-understand examples of how small law firms will be able to use some of the technologies presented in this guide to make their operations more efficient. We explain the capabilities of such tools in more detail in the next section 5, where we also hope to shed some light on some peculiarities of their internal workings. Most of the technologies presented below already work in practice in some way and in some countries (not necessarily for small law firms). Some will work only if appropriate data becomes available for that domain and jurisdiction, and there are also promising technologies that are not yet feasible in legal practice, but have already demonstrated their usefulness in research projects in law. Naturally, this is just an arbitrary list of possible technologies, without any intention to be comprehensive.

We have not excluded any solution for the sole reason that there will probably never be a market large enough for the product to survive in an “average” EU jurisdiction. You can see below that many of the features currently seen as innovations will require considerable investment and continuous maintenance by the law firm, in terms of both time and money. This includes the law firm with enough revenue to be able to pay not only its partners, fee earners and administrative employees, but also the consultants and IT suppliers necessary to prepare, implement and use such AI tools. Even if the proper IT tools are available at the present time, without such investment in implementation, most innovations will never take root on their own.

We refer to our protagonist just as “the Lawyer”. She lives in the not-too-distant future, and she is practising law in a fictional, non-English speaking member state of the European Union, and is a partner in a small law firm (say, composed of two partners and two employees).

### 6.2. Bilateral contract negotiations on a platform and recording time

The Lawyer has received a message that a new appointment has just been booked with a yet unknown client using the firm’s new MetalErg bot and appointment system. MetalErg is a (fictional) messaging application introduced just a year ago with tremendous growth figures thanks to its promises of ‘stronger than anything’ end-to-end encryption. The Lawyer quickly scans the agreed appointment time: it is still more than 6 hours away, which gives plenty of time and so is not urgent, and does not require an immediate interruption to the Lawyer’s work.

Before this new meeting, the Lawyer tries to finish analysing the draft sales contract of Client Alpha, a small business trying to license its very popular digital widgets to AgileLean Bank, which would like to use the widget for its latest marketing campaign. The draft which has been negotiated was based on the well-known template the Lawyer crafted some years ago for Alpha, which was customised by the sales people at Alpha who loaded the draft onto the AgileLean Bank’s contract negotiation platform

(which is a third-party platform, but paid for by the bank). Not surprisingly, a number of provisions in the contract are not in line with the procurement department's expectations. Luckily, this is a relatively small and flexible bank, so the items on their "list of expectations" as configured on the platform is manageable manually on their platform (a mere list of 75 items, of which 60 can be automatically verified to a certain degree). The Lawyer makes a note to Alpha that next time, before submitting drafts to AgileLean Bank, they should first send the draft to her instead, because she already has some experience with this bank and could have easily accommodated half of their contractual requirements in the initial draft submitted. The Lawyer logs onto the negotiation platform, the platform checks the e-ID of the Lawyer and also her power of attorney based on an attached attribute certificate, and lets the Lawyer review the discrepancies.

Because the bank has subscribed to the Enterprise Package of the negotiation platform, the platform has already undertaken an automated check for 60 items out of the 75, and has highlighted the problematic provisions that seem not to be acceptable. It has even provided alternative wording for half of the highlighted provisions. Thanks to the responsiveness of the Lawyer, the lawyers of the bank have not yet seen this draft. That means the Lawyer is able to make changes first in a way that best suits Alpha's interest, and label most of the items as resolved by either accepting the suggested changes (even if with considerable amendments) or by explaining why the automatic suggestion is not necessary at all and is already included somewhere in the agreement. For most of the explanations where legal arguments were needed, she has inserted convincing ECLI and ELI references to the government provided legal database. Being an experienced lawyer, she already has a large private repository of amendment texts for typical requirements and explanations – this repository has now been expanded by the non-subject matter specific requirement list of this particular bank.

Having submitted all the changes, the Lawyer instantly notifies Alpha's sales team off-channel (that is, in a separate message outside the contract negotiation platform) of any issues that they will need to answer, and with confidential suggestions on how to proceed with those changes. She also attaches the necessary screenshots, with scanned and OCR'd texts of the relevant version, which she created because the platform prohibited her from downloading the current draft. It would be best if the sales team is able to finalise the text before the bank's lawyers have time to react...

She quickly checks if her time was correctly recorded by her time tracking software. She has to consolidate the work on the message to Alpha and the work on the negotiating platform as a single activity. She also had to deduct the time spent on updating her private repository which is not billable. Otherwise, the automatic tracking was correct. Her message sent to Alpha has also been automatically filed to the correct matter and her process tracking information on ongoing work was also updated. Even if she is on a monthly fixed fee arrangement with Alpha, they still want to see how much work she actually does for them, to the point of details in minutes, otherwise they might fail to see her added value and try again to renegotiate their fee arrangement.

### 6.3. Client meeting and intake

The Lawyer now checks the details of the new appointment. The appointment was initiated through the MetalErg platform, but the details of the appointment were recorded by an app connected to the firm's own website (operating within the EU).

All she can see as detail is that the client has made a prepayment by a cryptocurrency for an hour of legal advice, and the client will expect the discussion to take place via the MetalErg platform. Due to privacy and deontology concerns, no other details can be recorded when using such a channel for an appointment (such as exactly what the questions will be), so there is no need for her to make preparations.

It is time for the new appointment and time to see (or hear or read material from) the new client. The MetalErg bot has already notified the prospective client — by way of pointing to the general terms and conditions of the law firm on the website — that due to deontology rules, it is not possible to provide

legal services other than oral advice via the MetalErg, and in which circumstances the lawyer will need to identify the client prior to providing any service.

The information the client needed was in relation to a cross-border succession law question of some very high-value Robed Ape non-fungible (NFT) tokens and half a pair of silver avatar gloves in Sandbox that he inherited from a late friend, but was not included in the European Certificate of Succession. Luckily, the Lawyer was well-versed in the issue, and so the client was soon satisfied after 40 minutes of intense question-answer dialogue. He was so impressed that he asked for another appointment regarding a court case he would like to initiate. Considering the remaining 20 minutes had been prepaid, the Lawyer suggested changing from instant messaging to a video channel so as to carry out the client identification and discuss some basics of the court case.

Considering that the client already had all the necessary information in his European digital identity wallet, the identification of Mr Beta took about 30 seconds.. A minute later all the necessary identification data was in the Lawyer's practice management system, confirmed and verified, with national know-your-customer rules complied with.

The court case is about unfair contractual terms in a consumer contract and related compensation based on a privacy breach, and Mr Beta would prefer a fixed fee arrangement. Based on this information, the Lawyer suggested that the next consultation should be two hours long. Based on that discussion, she will be able to give a fixed fee offer for drafting and filing. The Lawyer and Mr Beta enters into an engagement agreement for the two hour consultation to specify the details of the new case, and also agree by video on the date of the next consultation (the next day in the morning).

#### 6.4. A quick lease agreement is needed

Before the end of the workday, she wants to complete the request she has received two days ago for a new lease agreement from another client, Ms Gamma. Ms Gamma has a new prospective tenant for her apartment, and they have already agreed on the most commercially important parts (like the term of the lease, amount of monthly rent payable, the deposit, responsibility for repairs and maintenance etc.). Ms Gamma has already prepared a first draft she has used for the negotiation phase: the text came from the automated document assembly tool the law firm makes available to its clients (in exchange for a small monthly fee). As usual, no matter how sophisticated the template options may be, there are always unique requirements that Ms Gamma needs, and she has had enough bad experience to customise the text herself (and besides, the insurance company suggested that she do so).

Among other things, they have agreed with the tenant (a painter) that the tenant himself will pay for the renovations both at the start and at the end of the lease, and for any repainting needed during the lease, in exchange for a small deduction in the rent. Of course, Ms Gamma still wants to claim tax relief for these costs and is not sure how that will work. Although the changes required by the landlord are straightforward, the Lawyer reviews the full agreement, and makes some further adjustments due to regulatory changes effective from 1st January that are not yet incorporated into the template. She records her time on Ms Gamma's matter, sends a task to the trainee to include the necessary changes in all the relevant document assembly templates, and messages her law firm partner to discuss planned regulatory updates for the templates.

Ms Gamma has also asked the Lawyer to make a for-information translation in English of the final agreement for the new tenant as well. She sends the document for translation to the trainee, but of course, he will use an automated translation service for the first draft, and will only review the end result before sending the finished work to Ms Gamma.

## 6.5. Preparing for court work

During the two-hour consultation with Mr Beta, the Lawyer first carefully discusses with the client how and where he has bought the doorbell that is the subject of the dispute, and gathers the list of evidence she will need from the client: both the usual evidence (such as the webshop product page, the placing of the order, the information and terms and conditions published, the acknowledgement of the order received etc.) and the more specific evidence on the terms of the fixed subscription, including home insurance and monitoring, the evidence for faulty performance (when the owner was not recognised), the date of the data breach, the saved configuration of the doorbell, the breach notification sent by the provider, and the time of call with customer support when promises for free cancellation were made. At the end of the consultation, the Lawyer gives Mr Beta the list of evidence at his disposal to send to her, and she promises Mr Beta that she will come back to him with the offer within two further working days.

Based on the discussions, her memo and the evidence received so far, the Lawyer goes to the bidding module of her practice management system for some calculations. First, she enters the court case details such as the type of proceedings, the court to which the claim will be addressed, the causes of action and the list of evidence currently known. She now sees a statistic on the duration of the proceedings and the number of court hearings that were necessary in similar past cases. However, due to anonymisation concerns and the lack of a sufficient number of similar cases, there is not enough data to give her a reliable prediction on the amount of compensation that they may recover for the data breach and the leak of his video footage.

Considering that she does not yet know Mr Beta well, she will have to include a risk premium in her fixed fee calculation, as recommended by the bidding module. However, most of the missing evidence will be probably easy to acquire, because no cross-border requests are to be sent to providers other than for the audio recordings of the support call. Furthermore, Mr Beta sent all the required evidence in a very short time and in good quality, so she adjusts the evidentiary risk factor accordingly. She finalises the offer, and a draft engagement agreement is assembled that also includes the fees and terms of engagement.

## 6.6. A service for a professional entrepreneur

Next day, the Lawyer has to start the work she received last week from the EU network of independent firms to which her firm belongs: a country-level regulatory review for one of Ms Delta's new startups. Ms Delta is a successful serial entrepreneur from Nigeria. The new company is based on an integrated recycling platform of packaging of FMCG<sup>73</sup> products centred around ultra-low-cost flexible NFC<sup>74</sup> tags. The company sells and rents to both consumers and companies a huge array of popular recycling equipment integrated with NFC tags, and wishes to keep operations lean by using a decentralised application ("dApp") built from smart contracts and running on a permissioned blockchain. These contracts are each a piece of code, and together they ensure that customers no longer subscribing to certain features of the recycling equipment (such as accessing the waste exchange to sell recycled components) will no longer be able to use the equipment for these features.

Because Ms Delta already had a number of similar platforms, the company did not spend much time on creating and testing these new smart contracts, but still sent them for strict security audits. However, the insurance company said it will not cover the new platform with product liability insurance unless the company makes a country-by-country level risk review of how these smart contracts interact with and fit into the regulatory regimes of EU countries (a legal audit of the decentralised application). So, the company made an architectural overview of the platform, explaining how it is supposed to work, and all the smart contracts implemented in the dApp were also included in source code format (in

<sup>73</sup> FMCG: fast-moving consumer goods, see [here](#)

<sup>74</sup> NFC means near-field communications, NFC tags are tiny computers (system-on-a-chip) or memory devices that communicate with readers and writers via radio-frequency in short range, and enable precise tracking of goods.

the programming language Go). Based on offers submitted by multiple networks and law firms, the company selected the network of which the Lawyer's law firm is a member.

Of course, the sort of work required in this case is not very common among law firms. To start with, very few companies want to invest money in initial legal audits when they build a new dApp (and even fewer want to spend on periodic review of regulatory changes). Investors usually require review only for the largest jurisdictions, and even in those countries it is sufficient to involve only a handful of law firms to make the audits. The Lawyer is one of the very few lawyers in her country who has undertaken such work previously, and so she already has some experience in undertaking this review. She has also taken some training on this kind of work organised by the network.

Her job mainly involves familiarising herself with the architectural overview, not the code. She first adapts, in light of her own country's national laws, the template list of contentious situations, cases of typical government interventions and judicial orders which the dApp should be able to accommodate, and checks those against the process set out in the overview. Whether in code, by human administrative support or by special instructions of company officers, the startup should be able to address all items in the list. The lead network firm has a contract with a software specialist company which is able to compare the architectural overview with the source codes included, and if necessary, drill down in the latter and evaluate which situations are not covered by the smart contracts. Based on this internal review, the network firms will receive an updated version and finalise their first report to the startup.

## 6.7. Court work for Mr Beta and submission of the package to the court

Following Mr Beta's approval, she starts working on his case. She transfers existing evidence from the client communication files to the case management software, and records all other details she has. She checks how the statement of facts relates to the list of existing evidence, what evidence she will need to acquire prior to creating the first draft of the claim and how these pieces of evidence relate to the causes of action (compensation for privacy breach and refund for the early termination of the subscription).

Now she turns to the relevant case law. She starts researching in her legal database: what arguments and what evidence have courts accepted in the past from claimants as proof of data breaches related to consumer devices or home security and alarm services? She also investigates findings of unfair terms in consumer contracts in relation to non-performance of the supplier's obligations while obliging the consumer to fulfil the consumers' own obligations — in what previous cases has the unfairness of this condition been established, and what was the argumentation used by the claimants? She takes notes based only on her research.

A couple of weeks later, she now has all the evidence she needs to finalise the claim. She checks in the case management software for both causes of action that all the facts necessary to establish her claim are included, and all the facts are supported by evidence. Based on this, she reviews the draft statement of facts written by the software from the raw data, and rewrites this to create a more fluid, easy-to-read statement of facts. Now it is time for the software to generate the first draft of the claim while using the appropriate smart electronic form, including all the structured metadata for the pieces of evidence, her semi-structured reasoning and arguments marked-up with the correct computer representational language, and the non-structured narrative parts intended to convince the judges.

When she feels the document is ready, she creates the final package with the help of the case management software, including every attachment needed from the power of attorney to the last evidence and her electronic signature, and delivers it to the courts via a registered channel. The case management software automatically saves the acknowledgement of receipt. Within two minutes, the results of the validity and IT risk checks carried out on the package submitted and the new case identifier are received and recorded.



## 7. Risks to professional obligations when using AI tools

### 7.1. Introduction

The amount of information to be processed by lawyers continues to increase every year: a more digitised society creates more and more data. A surge in regulatory instruments created or court cases published boosts the demand for the processing capacities of lawyers. All this intensifies competition among legal service providers on the basis of the capabilities of their IT tools.

We have to remember that greater processing capacity for a lawyer does not necessarily mean either higher quality of work for the client, greater legal certainty or a more complete adherence to the rule of law. Similarly, it does not follow that lawyers will necessarily provide better services just because they have access to more regulatory instruments and cases.

When lawyers try to meet expectations with regard to the increase of their data processing capabilities, they need to ensure that they act in compliance with the core principles of the European legal profession. These principles are not only enshrined in deontological rules, but also in the jurisprudence of the European Courts and the European Court of Human Rights. These principles include especially the independence of the lawyer, the obligation to avoid any conflicts of interest and the duty to respect professional secrecy.

*When lawyers try to meet expectations with regard to the increase of their data processing capabilities, they need to ensure that they act in compliance with the core principles of the European legal profession.*

In this part, we highlight some risks to lawyers' compliance with their professional obligations when using AI tools. Some of these risks are more closely connected to the nature of the tools used or their delivery method (such as risks of using tools provided by third parties, risks of bias and lack of transparency in machine learning methods), and some risks arise from the perspective of the professional obligations of lawyers, such as the obligations of confidentiality or of competence.<sup>75</sup>

<sup>75</sup> See also e.g., Conseil National de Barreaux, Assemblée générale du 9 octobre 2020 Groupe de travail Legaltech (n 25) 14–17.

### 7.2.1. Risks arising from using cloud computing and online platforms to provide access to AI tools

The first risk of a technological nature that we require to mention is not specific to AI tools per se, but to the most popular way of deploying AI tools: using cloud services intended to be used by everyone (public cloud services).<sup>76</sup> The very same risks apply to lawyers using online platforms, because these platforms also rely on public cloud services,<sup>77</sup> and from this point of view, the source of the risk is the deployment model and not e.g. the double-sidedness of the model.<sup>78</sup>

In theory, any AI tools could work on-premise as well, so the way that AI tools are deployed should not be seen as an AI tool-specific risk. However, because of the severity of the effects of such risks and these risks affect and will affect almost all small law firms, it is indispensable that we discuss these risks – even if we know that these risks are not directly and necessarily tied to AI technology. We consider the risks of the deployment model as one of the most important risks of AI tools.

Cloud computing-based tools are popular for lawyers for the same reason that they are popular for all consumers: they avoid serious difficulties relating to implementation, configuration and maintenance, and keep technical tools both simple and cost-effective to operate. It becomes possible for lawyers to start using a complex service right away, with minimum upfront investment in time and money. They may also give the lawyer a promise of instant and secure access from anywhere.

The risks of cloud computing have been analysed in a greater deal of detail by the CCBE.<sup>79</sup> With regard to AI tools, the CCBE has highlighted issues such as the problem of extraterritoriality (users not having any control over how local regulations applicable to the cloud provider might affect rights and protections granted to lawyers in their home jurisdiction),<sup>80</sup> and the problem of how lawyers may access their own data at the end of the contractual relationship with the provider.<sup>81</sup> (The risks of privacy, governmental and unauthorised access are discussed in a separate section 7.2.3.)

The most important risk that which persist to this day unchanged is the risk of vendor lock-in. A minor risk could be that of resolving an eventual dispute with the service provider.<sup>82</sup> However, in our view, the bigger problem is the one described as “getting the data out of the cloud is much harder than putting it in”.<sup>83</sup> Separate documents, bills, balances, and client account information can be exported easily from a cloud computing service, but there is no standardised way for exporting all the transactional information and transferring it to another cloud service provider. The more completely the AI tool is integrated into the operations of a law firm, the bigger this problem becomes (especially for practice management related AI tools). The problem can be solved only in theory, by way of defining and using standards, but any such standardisation would need considerable investment specific to popular solutions available at the time of defining the standards. In practice, in an EU-country level fragmented market, this is very difficult to achieve.

This also means that the more successful a cloud provider becomes in terms of IT tools for lawyers, the more integrated its offering becomes (e.g., for practice management software providers), and the stronger the cloud service provider also becomes vis-à-vis lawyers.

76 Peter Mell and Tim Grance, ‘[The NIST Definition of Cloud Computing](#)’ (National Institute of Standards and Technology 2011) NIST Special Publication (SP) 800-145 7 accessed 5 December 2021.

77 Council of Bars and Law Societies of Europe, ‘[CCBE Guide on the Use of Online Legal Platforms](#)’ (29 June 2018) accessed 5 December 2021.

78 *ibid* 6.

79 Council of Bars and Law Societies of Europe, ‘[CCBE Guidelines on the Use of Cloud Computing Services by Lawyers](#)’ (7 September 2012) accessed 5 December 2021.

80 Council of Bars and Law Societies of Europe, ‘CCBE Guide on the Use of Online Legal Platforms’ (n 77) 11.

81 *ibid* 13.

82 *ibid*.

83 Homoki (n 5) 68.

*Getting the data out of the cloud is much harder than putting it in ... the more successful a cloud provider becomes in terms of IT tools for lawyers ... the stronger the cloud service provider becomes vis-à-vis lawyers*

### 7.2.2. Relying on results without proper explanation and understanding, and others risks relevant to the performance of AI tools

Based on the AI Index of 2021, after cybersecurity and regulatory relevance, the third most important risk in adopting AI tools that organisations consider relevant is the lack of explainability of the results of such tools.<sup>84</sup> The Commission has also highlighted that “in order to increase transparency and minimise the risk of bias or error, AI systems should be developed in a manner which allows humans to understand (the basis of) their actions”.<sup>85</sup>

One has to be aware of the problem that explainability has a very different meaning for technical people, for lay people and for court use. The Ethics Guidelines on Trustworthy AI sets out a requirement that “explanation should be timely and adapted to the expertise of the stakeholder concerned (e.g., layperson, regulator or researcher)”.<sup>86</sup>

Currently, this requirement on explainability is nothing more than a research objective that is very distant from reality.

Already a large number of research papers have been published that deal with explainability specifically within the domain of the law,<sup>87</sup> but even the most basic concepts of explainability and the classifications of types of explanations are diverse, and are also domain- and use-case specific.<sup>88</sup>

This highlights a very significant problem. Due to their increased performance, NLP techniques such as word embedding language models are becoming dominant in the legal use of AI.<sup>89</sup> However, this technique is a black box type of model, meaning that these models do not operate in a way that can be interpreted and explained.<sup>90</sup> Evidently, however, legal uses and the principle of rule of law demand a strong-sense interpretability of the results of a model.<sup>91</sup>

[box: Word embedding language models are becoming dominant in the legal use of AI. However, this technique is a black box type of model, meaning that these models do not operate in a way that can be interpreted and explained. Evidently, however, legal uses and the principle of rule of law demand a strong-sense interpretability of the results of a model.]

When an AI tool is used in a low-risk context, such as when the results are expected to be reviewed by a professional, prior to being sent to clients, explainability tends not to be so much of a problem. In such a case, the professional is also expected to check not only the validity of the results but also if the relevant facts of the case provide adequate support to explain the results. But such a requirement of a step-by-step human review severely limits the usability of AI tools.

Even if only humans have the true capability to understand a written text in versatile ways, such a review is more prone to fatigue, stress and emotion and it is usually considerable slower. These factors all decrease human performance compared to the models used by AI tools. That means the longer the results of the AI tool become (i.e. the longer the generated text becomes, the longer the list of search results is, the more text that needs to be categorised and labelled), the less we can rely on human review mitigating risks from the AI tool, including the problem of the lack of explainability of the results.

84 Daniel Zhang and others, ‘[The AI Index 2021 Annual Report](#)’ [2021] arXiv:2103.06312 [cs] 102 accessed 4 December 2021.

85 European Commission (n 2) 14.

86 High-Level Expert Group on Artificial Intelligence, ‘[Ethics Guidelines for Trustworthy AI](#)’ accessed 12 December 2021. 20.

87 Giulia Vilone and Luca Longo, ‘[Explainable Artificial Intelligence: A Systematic Review](#)’ [2020] arXiv:2006.00093 [cs] 6 accessed 12 December 2021.

88 *ibid* 7–14.

89 Vadász and others (n 6).

90 Prashant Gohel, Priyanka Singh and Manoranjan Mohanty, ‘[Explainable AI: Current Status and Future Directions](#)’ [2021] arXiv:2107.07045 [cs] 11 accessed 12 December 2021.

91 Council of Bars and Law Societies of Europe, ‘CCBE Considerations on the Legal Aspects of AI’ (n 3) 12.

*That means the longer the results of the AI tool become, the less we can rely on human review mitigating risks from the AI tool, including the problem of the lack of explainability of the results.*

Requiring step-by-step human review limits the number of AI tools which can be interconnected with each other: even the basic AI tools available today often rely on many different layers of automated decisions. For instance, for a legal due diligence review tool, documents are first labelled (categorised) by their languages, and then by document types (e.g. based on three different, non-overlapping taxonomies), then labelled based on the document types, and finally, relevant information is extracted from the documents (such as provisions on term and termination, amount of rent etc.). In the next step, a risk score is estimated by the AI tool for that given document and finally, the lawyer receives an automatically generated report on all the documents reviewed. Integrating a compulsory human review at every step would make this tool too difficult to use, which makes this approach both impractical and unlikely.

Publishers of AI tools tend to provide as little transparency on the internal working of their AI tool as is they can get away with whilst seeking to convince lawyers to subscribe to their tools. This is understandable, because by explaining in detail how their successful tools work (including how the dataset was built, what architecture they are using etc.), they will just face more intense competition by free-riders closely following their approach. Current intellectual property rights cannot fully address the concerns of publishers of AI tools, and so they try to minimise the information needed for transparency, and so also for explainability.

That also means that requirements for transparency and explainability have to be demanded by lawyers themselves, the users of such tools, who should not just simply leave it to the publishers to fulfil this general social need.

Besides explainability and transparency, there is a number of other characteristic problems arising from AI tools.<sup>92</sup> One well-known such problem is what is called the brittleness of AI: the system functions well within certain bounds but poorly outside those bounds.<sup>93</sup> That is, it may seem to perform well as long as the real-life conditions are similar to the training conditions, but the performance might be unexpectedly and severely degraded in certain situations not encountered during the original training. Changes that are minor or even invisible for human attention (such as putting some small extra stickers on a STOP traffic sign) might cause the AI to misinterpret a STOP sign as a speed limit sign.<sup>94</sup> This brittleness of an AI can indicate deficiencies in both the reliability and training of AI tools. Such unexpected results may show that the tool is simply not yet reliable in the live operating conditions of its intended use, and this lack of reliability might be improved by providing more diverse training data for the tool. Legal use of AI tools is less critical in terms of dangers to life and limb compared to well-known AI uses such as self-driving cars, automated weapons and avionics,<sup>95</sup> but some standards developed in the validation of AI tools in the latter domains can also be used with legal AI as well.

Another characteristic problem is only partly of a technical nature: bias and discrimination or unfairness may be caused both by the data source used and by the algorithms and models evaluating the data.<sup>96</sup> Bias may result from historical data skewed towards specific groups or from specific data collection mechanics, from lack of data in certain areas<sup>97</sup> (also called unobservable outcomes or survivorship

92 See also Charles Q. Choi, '[7 Revealing Ways AIs Fail](#)' (*IEEE Spectrum*, 21 September 2021) accessed 14 December 2021. This list of characteristics problem is just an illustration for lawyers, and does not intend to provide any comprehensive view on such subject, omitting important other problems like catastrophic forgetting etc.

93 Andrew J. Lohn, '[Estimating the Brittleness of AI: Safety Integrity Levels and the Need for Testing Out-Of-Distribution Performance](#)' [2020] arXiv:2009.00802 [cs, stat] 1–2 accessed 14 December 2021.

94 Douglas Heaven, '[Why Deep-Learning AIs Are so Easy to Fool](#)' (2019) 574 *Nature* 163 accessed 25 February 2022.

95 Lohn (n 93) 5–6.

96 For examples in the legal field, see Council of Bars and Law Societies of Europe, 'CCBE Considerations on the Legal Aspects of AI' (n 3) 23–24.

97 Nengfeng Zhou and others, '[Bias, Fairness, and Accountability with AI and ML Algorithms](#)' [2021] arXiv:2105.06558 [cs, stat] 5 accessed 18 December 2021.

bias<sup>98</sup>), both resulting in what is called training data bias.<sup>99</sup> Equally important is that bias can also be introduced by reason of the chosen methods, models and architects, including statistical rules or the optimisation methods and parameters for performance.<sup>100</sup> A third major source of bias is related to a change in the usage of the AI tool: systems originally intended for one purpose may be tried for reuse in a different area (including cross-jurisdictional use of AI, e.g. a legal analytics tool popular in the USA used by lawyers in Ireland) or the results of a tool intended for a different context are simply misinterpreted by a user (or another AI tool) working in a slightly different context.<sup>101</sup>

### 7.2.3. Risks to privacy

Most, but not all of the privacy risks are related to AI tools deployed in cloud services. In the last nine years, cloud computing services have matured significantly, and have become mainstream even in the most tightly regulated sectors, such as in financial regulation.<sup>102</sup> Security processes have become more robust with widely accepted third-party attestations and assurances on the reliability of IT security controls. This includes the Cloud Control Matrix<sup>103</sup> or most of the reports of the Service Organization Controls of the American Institute of Certified Public Accountants.<sup>104</sup> However, certain information security controls are still lacking. As long as the service provider (or any underlying platform or infrastructure provider) is technically able to read and access the data of the lawyer, the risks of unauthorised access will remain a serious concern for lawyers. We discuss this in more detail in relation to the professional obligation of confidentiality, in section 7.3.3, except for one aspect, the danger of an online service provider reusing client data for its own use (including for new trained models), which we discuss here.

Regarding reuse of data, the advice given in relation to online platforms applies here as well.<sup>105</sup> The most striking problem is that current online legal platforms, including those providing AI tools, rarely provide any information at all about the possibility for the provider to reuse information at the provider's disposal. Lawyers should seek to insist upon terms and conditions clearly excluding any profiling activity (even if the target for the profiling is only the same lawyer) and reuse of data even after supposed anonymisation of the data.

*Lawyers should seek to insist upon on terms and conditions clearly excluding any profiling activity (even if the target for the profiling is only the same lawyer) and reuse of data even after supposed anonymisation of the data.*

The problem is that many forms of unstructured legal data – such as full text judicial decisions or contracts – are very hard to anonymise, and it is difficult to remove personal data from them. Simply removing names and locations is not enough, because a number of unique events or other context-related information may be present in a decision which make it possible to reidentify the client or even narrow down the number of possible relevant entities, which in itself would result in a breach of confidentiality.<sup>106</sup>

98 'Survivorship Bias', *Wikipedia* (2021) accessed 18 December 2021.

99 Xavier Ferrer and others, 'Bias and Discrimination in AI: A Cross-Disciplinary Perspective' (2021) 40 *IEEE Technology and Society Magazine* 72, 1.

100 David Danks and Alex John London, 'Algorithmic Bias in Autonomous Systems', *Electronic proceedings of IJCAI 2017* (2017) 4693 accessed 18 December 2021.

101 *ibid* 4694.

102 See e.g., 'Proposal for a Regulation of the European Parliament and of the Council on Digital Operational Resilience for the Financial Sector and Amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014 and (EU) No 909/2014' accessed 12 December 2021. or European Banking Authority, 'Guidelines on Outsourcing Arrangements' (5 June 2019) accessed 12 December 2021.

103 Cloud Security Alliance, see <https://cloudsecurityalliance.org/research/cloud-controls-matrix/>

104 See also Council of Bars and Law Societies of Europe, 'CCBE Guide on the Use of Online Legal Platforms' (n 77) 11.

105 *ibid* 12.

106 Gergely Márk Csányi and others, 'Challenges and Open Problems of Legal Document Anonymization' (2021) 13 *Symmetry* 1490, accessed 25 February 2022.

So, this risk persists even if a service provider claims that it will anonymise all stored data before using it for any purposes other than service to the lawyer (e.g., before reselling). As mentioned in the CCBE Guide on online platforms,<sup>107</sup> in spite of all best endeavours, it might turn out that the anonymised dataset can later be reidentified, based on further information that is at the disposal of a third person who may have gained access to the originally anonymised dataset, and safeguards in GDPR (such as Article 22) may also not be able to resolve this problem.

Beside the problems common to AI tools, cloud services and online platforms, it is also important to highlight a specific privacy risk that is applicable to AI tools only. Even when undertaking processing on-site (i.e. not using any cloud computing services), a lawyer could be in breach of confidentiality obligations if the law firm itself trains a model based on client data in its possession (such as timesheet information, contracts, their own court documents), and then allows a third party to copy this trained model data and take it for its own use. The output of training (trained models) often does not appear to be legible data at all, but one should be mindful of the problem that the output of popular language embeddings (texts turned to numerical representations, like vectors of the text, such as in BERT etc.) can be used to predict the original text used for (pre)training, and reverse-engineer it to disclose sensitive information from the source.<sup>108</sup>

*The most striking problem is that current online legal platforms, including those providing AI tools, rarely provide any information at all about the possibility for the provider to reuse information at the provider's disposal ... in spite of all best endeavours, it might turn out that the anonymised dataset can later be reidentified, based on further information that is at the disposal of a third person ... One should be mindful of the problem that the output of popular language embeddings can be used to predict the original text used for (pre)training.*

This also demonstrates why lawyers should never authorise AI tool providers to use legal data at their disposal for training or analysis by third parties unless the lawyer obtains reliable assurances that the given method does not suffer from similar privacy risks. Without such an assurance, the lawyer is simply not in a position to use such a service. Simply accepting standard terms and conditions containing implicit or explicit permissions to train or analyse would probably be in violation of professional obligations of confidentiality unless the client has provided a prior informed consent. However, such consent is not likely to be given if the lawyer is not even aware of the risks.

<sup>107</sup> Council of Bars and Law Societies of Europe, 'CCBE Guide on the Use of Online Legal Platforms' (n 77).

<sup>108</sup> Xudong Pan and others, '[Privacy Risks of General-Purpose Language Models](#)', 2020 IEEE Symposium on Security and Privacy (SP) (2020) accessed 18 December 2021.

## 7.3. Risks arising from professional obligations

### 7.3.1. Risks to professional competence: dangers of trying out new technology

Changes in technology may affect what questions are asked of lawyers and what answers lawyers should give (the knowledge base a lawyer works with), but these may also alter how lawyers operate their legal practices.

Such changes may create a conflict between the expectations of some clients and the professional expectation of a competent approach by the lawyer. Clients' wishes and business development reasons could both create an incentive for lawyers to appear as a pioneer in trying out (and relying on) new technologies and also to «boldly give advice» in areas where no lawyer has done so before.

As for the first aspect, lawyers are often expected to guess how society and courts will react to the impacts of a technical tool. Clients can exhibit overconfidence and shallow knowledge of a technology, regardless of whether they are inventors, professional investors or just private persons in the know. However, the professional expectation of competence for lawyers requires that they should not “take on a case which he or she is not competent to deal with”.<sup>109</sup> Lawyers must act in accordance with this principle at all times, even if a client suggests relying instead on a technological tool. Continuing professional development for lawyers is based on the appropriate training of the lawyer, which in turn implies some previous accumulation of experience that is used for that training. The proper focus of a lawyer's field of work is the social impact of technology, and not on understanding how technology works and, based on that, to make predictions on what future impacts a piece of technology will have on society. Of course, there are many ways lawyers can provide appropriate advice to their clients regarding the results of innovative technologies, but sometimes the best advice a lawyer can give is to communicate clearly any uncertainty and to demonstrate how weak predictions may be, and how misleading past forecasts are.

With regard to the internal working of a law firm, like in any company, the fear of missing out (FOMO) in the field of AI may increase risks related to insufficient professional competence. The amount of information annually published on AI is enormous even for academics (e.g. in 2020, this increased by 34.5% from the previous year<sup>110</sup>), and so coping with so much new information is difficult even for specialists in this technology. Publishers of AI tools occasionally provide incomplete or even misleading information about their products, make products available that are not yet suitably tested, or not sufficiently customised for the specific needs of a lawyer in a given jurisdiction, or they simply disregard applicable deontological or other rules.

Altogether, the feeling of FOMO and the lack of reliable information may lead lawyers to make bad decisions which are in contrast to the professional obligation of competence. They may invest in unreliable technology and build important processes of their law practice on tools that have still not been properly tested on the given (national) market, or are not yet in line with the national rules of deontology. The major risk is not in spending the income of the practice unwisely – that is the least of the problems. An over-eagerness in trying out new tools may lead to the more strategic problems of unexpected data breaches or violation of professional obligations.

*With regard to the internal working of a law firm, like in any company, the fear of missing out (FOMO) in the field of AI may increase risks related to insufficient professional competence.*

<sup>109</sup> Council of Bars and Law Societies of Europe, ‘[Charter of Core Principles of the European Legal Profession and Code of Conduct for European Lawyers](#)’ (2021) 8 accessed 15 March 2022.

<sup>110</sup> Zhang D and others, ‘[The AI Index 2021 Annual Report](#)’ [2021] arXiv:2103.06312 [cs] 25 accessed 4 December 2021

Lawyers are expected to be aware of their deficiencies as part of the obligation of competence. It is important for small law firms to identify what processes need to be in place at the law firm to be able to use new, promising tools reliably. But publishers of tools sometimes promise simple technological means as a substitute for non-existing manual processes in firms, which is not a prudent approach from the viewpoint of the user (the publishers of these tools may be unaware of how the manual process would have worked, why process steps were important etc.). A technical tool substituting or incorporating a process will not provide the same flexibility as that of which a manual solution would be capable.

### **7.3.2. Risks to professional competence: integrating technical and human processes, balancing promises with actual capabilities**

The expectation of the competence of a lawyer also demands that lawyers do not take on more clients than they are able to serve professionally. This is a considerable risk when lawyers have cheap access to very effective tools which sometimes provide unexpected increases in visibility, made possible by a more interconnected, digital society.

AI tools can easily increase this risk. These tools can significantly contribute to making law firms more cost effective, e.g. by making client intake more streamlined, carrying out legal research and document creation faster and without any repetition. But there is always a strict limit to the maximum volume of service that a single law firm can undertake while working according to the same business model (remaining a law firm).

Activities in law firms are built around the responsibilities of highly skilled individuals. These professionals and their knowledge serve as the primary cornerstone of the law firm's business. This is markedly different in other businesses, where the most valuable parts of a business (the core competitive advantage) may be a well-developed sales channel, complex business processes implemented, the firm's equity or other impersonal capital assets of the business.

A major reason for this personal nature of law firms is the requirement to understand what a client really needs in terms of legal advice compared to what the client is able to articulate as a legal need. Even if a lawyer receives a large number of very similar requests from consumers or micro enterprises, the evaluation of clients' actual needs requires the attention of a trained, professional lawyer.

*A major reason for this personal nature of law firms is the requirement to understand what a client really needs in terms of legal advice compared to what the client is able to articulate as a legal need. Even if a lawyer receives a large number of very similar requests from consumers or micro enterprises, the evaluation of the clients' actual needs requires the attention of a trained, professional lawyer.*

Therefore, an advanced, fully automated operational structure of a hypothetical future small law firm will not be able to provide more services than a certain number of hours a day per trained lawyer, no matter the demand from society for such services. Otherwise, the business model is no longer that of a law firm.

Because of the core principles of the legal profession, lawyers are expected to avoid taking on more client engagements than they are able to serve appropriately. Even if, deontologically, lawyers are free to refuse requests from a large number of potential clients (due to lack of capacity), when relying on AI tools to attract and engage with new customers it might be difficult to draw a clear line between rejecting a potential client and failing to serve a client already engaged. Due to the special characteristics of the profession of lawyers, there are countries and situations where it is not an option simply to let clients

down in such cases, and to let them find a new lawyer if they want to. Overcommitment by individual law firms can often lead to economic problems and insolvency. In a number of countries, there are specific regulations on the liquidation of law firms. In such situations, the infringement of the obligation to act within one's professional competence leads to capacity problems and overcommitments, and becomes the problem both of other licensed law firms and of bars.

### **7.3.3. Risks related to professional secrecy obligations of the lawyer**

Many aspects discussed in section 7.2.1 and 7.2.3 are also relevant to the professional secrecy obligation of lawyers, and are not repeated here.

As discussed above, information security controls are still lacking in cloud services where providers are often technically able to read and access the data of the lawyer, resulting in a risk of unauthorised access, and thus breaches of confidentiality and legal professional privilege obligations, including reuse of client data for other purposes<sup>111</sup> or unlawful interception of communications by authorities.

From the perspective of their deontological obligations, lawyers have to be aware of the privacy threats to which client data becomes exposed when lawyers choose certain data processors (including cloud service providers).

Evaluating client confidentiality is never merely a technical exercise, reviewing boring lists of IT security controls, or ISO 27000 or SOC reports. The most important question a lawyer has to take into account is not the costs or the ease of use, but the risks to the client in case of third-party access: who might be interested in the details, what could happen to the client in case of a successful attack, and where the lawyer can keep relevant records safe during the engagement.

Whenever third-party service providers are involved (whether cloud computing or not), a risk exists that a prosecutor or a judge or others may order a service provider to provide access to information that is protected by lawyer-client confidentiality, and at the same time, prohibit the provider from informing the lawyer of such access.<sup>112</sup> Also, despite clear statutory or other legal protections of client-lawyer confidentiality, or human right requirements for a fair trial, procedural means to enforce these requirements are still lacking in many situations, such as service providers not being aware that a protective regime should be in place vis-à-vis demands from law enforcement agencies due to the confidential nature of the information which is being processed. That might result in a case where electronic information located at the premises of the lawyer is better protected from access compared to the capabilities of a top-notch security cloud provider or data-centre. Of course, in recent decades, self-hosting basic services like e-mail servers reliably is becoming more difficult and expensive for law firms and other enterprises (continuous upgrades and more sophisticated configuration, maintenance of blacklists and spam filters, with clients expecting compliance with expensive measures like disaster recovery and testing).

In addition, although the regimes regulating government access to data at service providers within the EU are expected to become more harmonised in the future, these are still very far from being completely harmonised even within the EU, and mostly remain the subject of national regulation. Therefore, what might be lawful access in the country of the provider of an AI tool might not be lawful in the country of operation of the lawyer, and such differences jeopardise clients' interests.

### **7.3.4. Risks related to the independence of the lawyer**

Lawyers are expected to be free and independent in pursuing their activities in advising and representing the client. This includes independence from the state and any other entities, and lawyers must not allow their independence to be compromised by improper pressure for business reasons.

Similar to the risks of vendor lock-in already mentioned in section 7.2.1, lawyers should be aware of the risk that, in the longer term, the most successful AI tools may also negatively affect their independence. If lawyers can use a single tool for purposes that are important for their business processes, lawyers

<sup>111</sup> Council of Bars and Law Societies of Europe, 'CCBE Guide on the Use of Online Legal Platforms' (n 77) 12.

<sup>112</sup> Regarding this subject, please see Council of Bars and Law Societies of Europe, '[CCBE Comparative Study on Governmental Surveillance of Lawyers](#)' Data in the Cloud' (2014) accessed 7 February 2022.

will face the same problems as set out in the CCBE Guide on the Use of Online Legal Platforms, that is interference by the provider of the AI tool in the lawyers' relationship with clients (including that the AI tool provider might become a platform, even without trying expressly to sell its services as such). This could be a more acute problem in oligopolistic or monopoly markets, where bottleneck inputs could hinder new entrants from providing substitute services. Such bottleneck inputs might be the models (like foundation models<sup>113</sup>) or limited access to the necessary data sources. Besides bottleneck inputs, fragmented market sizes in small jurisdictions also increase the possibility of a single player dominating an important market in AI tools.<sup>114</sup> The economic weight of a widely used AI solution can force lawyers to accept terms set out by AI providers that infringe their independence. Lawyers, and also future regulators of innovative tools, need to be actively aware of this risk.<sup>115</sup>

---

113 Rishi Bommasani and others, '[On the Opportunities and Risks of Foundation Models](#)' [2021] arXiv:2108.07258 [cs] accessed 19 December 2021.

114 Homoki (n 5) 46–47.

115 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 (n 54).



## 8. Conclusions

This guide has attempted to provide a glimpse into the exciting opportunities that AI tools may hold for the legal profession, with a focus on how small law firms may benefit, and the risks of which they should be mindful.

Today, it can already be said that AI is not in a position to make the legal profession redundant, and that its tools will not set lawyers on a path similar to that of elevator operators. However, depending on how society will be transformed by the use and reliance of such technology, lawyers will also have to adapt and adjust their workflows and competences.

The transformation will not be without problems, and we should not forget the dangers to human rights, fair trials and our core principles. To prepare to defend these values better, lawyers, bars and law societies should understand the risks of transformation in more depth, and that can only be achieved by using AI tools with prudence.

This has to be done even if the current tools are far from perfect, or are adapted to different jurisdictions or languages, exacerbated by the fragmented market for IT tools for lawyers across the Member States of the EU. Current tools may not provide all the information deemed necessary by lawyers, and may lack financial sense or practical usability. But there are already AI applications that can be tried even by the smallest of law firms – of course, with sufficient caution, in specific low-risk areas, and all the while understanding the relevance of the risks set out in this guide.

The use of more and more AI tools is inevitable in the long run, because it is not something that any profession can avoid – nevertheless, lawyers should pay attention to be sure that such tools are used in a way that does not harm clients or the rule of law, and that our use of the tools does not in any way lead to reducing the protection for the weak that a legal system is expected to provide.

Finally, we should warn that this guide evaluates only certain areas of risk related to AI. The technologies mentioned here might still not be desirable with regard to their implications, for example, for the rule of law, the core values of the profession or human rights.



## 9. Bibliography

[‘Article 33 - LOI N° 2019-222 Du 23 Mars 2019 de Programmation 2018-2022 et de Réforme Pour La Justice - Légifrance’](#) accessed 28 December 2021

Ashley KD, [Artificial Intelligence and Legal Analytics](#) (First, Cambridge University Press 2017)

Bommasani R and others, [‘On the Opportunities and Risks of Foundation Models’](#) [2021] arXiv:2108.07258 [cs] accessed 19 December 2021

Bourne CP and Hahn TB, [A History of Online Information Services, 1963-1976](#) (Cambridge, Mass : MIT Press 2003) accessed 27 December 2021

Chalkidis I and others, [‘Neural Contract Element Extraction Revisited’](#) (2021) abs/2101.04355 CoRR

Choi CQ, [‘7 Revealing Ways AIs Fail’](#) (IEEE Spectrum, 21 September 2021) accessed 14 December 2021

Debra Cassens Weiss, [“‘Treated like a Robot,” Contract Lawyers Chafe under Fickle Facial Recognition Surveillance’](#) (ABA Journal, 15 November 2021) accessed 19 December 2021.

Conseil National de Barreaux, Assemblée générale du 9 octobre 2020 Groupe de travail Legaltech, [‘Legaltechs du domaine de la jurimétrie: Préconisations d’actions Rapport’](#) (9 October 2020) accessed 27 December 2021

Council of Bars and Law Societies of Europe, [‘CCBE Guidelines on the Use of Cloud Computing Services by Lawyers’](#) (7 September 2012) accessed 5 December 2021

Council of Bars and Law Societies of Europe, [‘CCBE Comparative Study on Governmental Surveillance of Lawyers’ Data in the Cloud’](#) (2014) accessed 7 February 2022

Council of Bars and Law Societies of Europe, [‘Charter of Core Principles of the European Legal Profession and Code of Conduct for European Lawyers’](#) (2021) accessed 15 March 2022

Council of Bars and Law Societies of Europe, [‘CCBE Guide on the Use of Online Legal Platforms’](#) (29 June 2018) accessed 5 December 2021

Council of Bars and Law Societies of Europe, [‘CCBE Considerations on the Legal Aspects of AI’](#) (2020) accessed 19 November 2021

Csányi GM and others, [‘Challenges and Open Problems of Legal Document Anonymization’](#) (2021) 13 Symmetry 1490 accessed 25 February 2022

David Danks and Alex John London, '[Algorithmic Bias in Autonomous Systems](#)', Electronic proceedings of IJCAI 2017 (2017) 4693 accessed 18 December 2021.

European Banking Authority, '[Guidelines on Outsourcing Arrangements](#)' (5 June 2019) accessed 12 December 2021

European Commission, '[Communication from the Commission: Artificial Intelligence for Europe](#)' (2018) accessed 19 November 2021

European Data Protection Board, '[Guidelines 02/2021 on Virtual Voice Assistants](#)' (7 July 2021) accessed 14 January 2022

High-Level Expert Group on Artificial Intelligence, '[Ethics Guidelines for Trustworthy AI](#)' accessed 12 December 2021

European Commission for the Efficiency of Justice (CEPEJ), '[Guidelines on Electronic Court Filing \(e-Filing\) and Digitalisation of Courts](#)' (9 December 2021) accessed 27 December 2021

Ferrer X and others, '[Bias and Discrimination in AI: A Cross-Disciplinary Perspective](#)' (2021) 40 IEEE Technology and Society Magazine 72 accessed 25 February 2022

Gohel P, Singh P and Mohanty M, '[Explainable AI: Current Status and Future Directions](#)' [2021] arXiv:2107.07045 [cs] accessed 12 December 2021

Heaven D, '[Why Deep-Learning AIs Are so Easy to Fool](#)' (2019) 574 Nature 163 accessed 25 February 2022

Homoki P, '[Overview on Average State of the Art IT Capabilities and Comparison with Best Practices United Kingdom, USA and Canada](#)' (Council of European Bars and Law Societies (CCBE), European Lawyers Foundation)

Lippi M and others, '[CLAUDETTE: An Automated Detector of Potentially Unfair Clauses in Online Terms of Service](#)' (2019) 27 Artificial Intelligence and Law 117 accessed 25 February 2022

Lippi M and Torroni P, '[Argumentation Mining: State of the Art and Emerging Trends](#)' (2016) 16 ACM Transactions on Internet Technology 1 accessed 25 February 2022

Lohn AJ, '[Estimating the Brittleness of AI: Safety Integrity Levels and the Need for Testing Out-Of-Distribution Performance](#)' [2020] arXiv:2009.00802 [cs, stat] accessed 14 December 2021

'[Managed by Bots: Surveillance of Gig Economy Workers](#)' (Privacy International) accessed 19 December 2021

Medvedeva M and others, '[Automatic Judgement Forecasting for Pending Applications of the European Court of Human Rights](#)' (2021) accessed 25 February 2022

Mell P and Grance T, '[The NIST Definition of Cloud Computing](#)' (National Institute of Standards and Technology 2011) NIST Special Publication (SP) 800-145 accessed 5 December 2021

Pasquale F, New Laws of Robotics: Defending Human Expertise in the Age of AI (The Belknap Press of Harvard University Press 2020)

Poudyal P and others, '[ECHR: Legal Corpus for Argument Mining](#)', ARGMINING (2020) accessed 25 February 2022

[‘Proposal for a Regulation of the European Parliament and of the Council Amending Regulation \(EU\) No 910/2014 as Regards Establishing a Framework for a European Digital Identity’](#) accessed 21 November 2021

[‘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’](#) accessed 12 December 2021

[‘Proposal for a Regulation of the European Parliament and of the Council on Digital Operational Resilience for the Financial Sector and Amending Regulations \(EC\) No 1060/2009, \(EU\) No 648/2012, \(EU\) No 600/2014 and \(EU\) No 909/2014’](#) accessed 12 December 2021

[‘Proposal for a Regulation of the European Parliament and of the Council on European Production and Preservation Orders for Electronic Evidence in Criminal Matters’](#) accessed 5 December 2021

[‘Survivorship Bias’](#), Wikipedia (2021) accessed 18 December 2021

Themis Solutions Inc., [‘2021 Legal Trends Report Published by Clio’](#) (August 2021) accessed 30 December 2021

Themis Solutions Inc., [‘Legal Trends Report 2017 Powered By Clio’](#) (2017) accessed 30 December 2021

Themis Solutions Inc., [‘Legal Trends Report 2018 Powered By Clio’](#) (2018) accessed 30 December 2021

Tippett EC and others, [‘Does Lawyering Matter? Predicting Judicial Decisions from Legal Briefs, and What That Means for Access to Justice’](#) [2021] Texas Law Review

Tuggener D and others, [‘LEDGAR: A Large-Scale Multi-Label Corpus for Text Classification of Legal Provisions in Contracts’](#), Proceedings of the 12th Language Resources and Evaluation Conference (European Language Resources Association 2020) accessed 10 October 2021

Vadász P and others, [‘A Report on the Barriers and Opportunities in the Use of Natural Language Processing Tools in Small Legal Offices’](#) (Council of European Bars and Law Societies, European Lawyers Foundation) accessed 25 February 2022

Vilone G and Longo L, [‘Explainable Artificial Intelligence: A Systematic Review’](#) [2020] arXiv:2006.00093 [cs] accessed 12 December 2021

Xudong Pan and others, [‘Privacy Risks of General-Purpose Language Models’](#), 2020 IEEE Symposium on Security and Privacy (SP) (2020) accessed 18 December 2021.

Zhang D and others, [‘The AI Index 2021 Annual Report’](#) [2021] arXiv:2103.06312 [cs] accessed 4 December 2021

Zhou N and others, [‘Bias, Fairness, and Accountability with AI and ML Algorithms’](#) [2021] arXiv:2105.06558 [cs, stat] accessed 18 December 2021.