




ARTICLE

Generative AI systems in legal practice offering quality legal services while upholding legal ethics

Kalliopi Terzidou 

Faculty of Law, Economics and Finance, Department of Law, University of Luxembourg, Luxembourg, Luxembourg
Email: kalliopi.terzidou@uni.lu

(Received 21 August 2024; revised 24 December 2024; accepted 2 February 2025)

Abstract

Generative artificial intelligence (AI) systems, notably ChatGPT, have emerged in legal practice, facilitating the completion of tasks, ranging from electronic communications to the drafting of documents. The generative capabilities of these systems underscore the duty of lawyers to competently represent their clients by keeping abreast of technological developments that can enhance the efficiency and effectiveness of their work. At the same time, the processing of clients' information through generative AI systems threatens to compromise their confidentiality if disclosed to third parties, including the systems' providers. The present paper aims to determine the impact of the use of generative AI systems by lawyers on the duties of competence and confidentiality. The findings derive from the application of doctrinal and empirical research on the legal practice and its digitalisation in Luxembourg. The paper finally reflects on the integration of generative AI systems in legal practice to raise the quality of legal services for clients.

Keywords: ChatGPT; client-centricity; competence; confidentiality; generative AI; large language models

1 Introduction

Lawyers have been using cloud-based services ever since their appearance in the early 2000s to facilitate their daily administrative tasks, ranging from electronic communications and the drafting of documents to the use of office applications over the Internet, on-demand and remotely (IBM 2024). In this context, generative AI systems stand out because they allow the generation of text by being trained on a large amount of data, which can reach up to trillions of data in the case of Large Language Models (LLMs), to understand natural-language inputs and perform an array of related tasks (IBM 2023). Besides chatbots, private companies offer LLMs as plug-ins to office applications that can suggest an initial template for documents, such as contracts, or prepare presentation slides by adding relevant content from a text source, such as a legal report (Spataro 2023). LLMs are also able to automate the translation of documents, not only for research but also for communication purposes.¹ Automated summarisation of long e-mail threads and of transcriptions of virtual meetings can also be of benefit to the efficiency of the daily work of lawyers.

ChatGPT is the fastest-growing consumer application in history (Hu 2023) and possibly the most popular general-purpose AI (GPAI) system to date. ChatGPT is characterised by 'the generality and the capability to competently perform a wide range of distinct tasks ... typically

¹In Luxembourg, the multiple official languages of the country – that is, Luxembourgish, German and French – along with other working languages corresponding to an international clientele, including English, render the translation of documents an everyday necessity for lawyers.

trained on large amounts of data, through various methods'.² Large generative AI models are typical examples of GPAIs, enabling the 'flexible generation of content, such as in the form of text, audio, images or video, that can readily accommodate a wide range of distinctive tasks'.³ A notorious use of ChatGPT in legal practice involved two lawyers from New York who prompted ChatGPT to retrieve case-law to support their client's arguments in the brief submitted to the court (Merken 2023). However, the six cases that the system suggested were fictitious, which led to the sanctioning of the lawyers by the US District Judge. This is not an uncommon incident, since ChatGPT is known for 'hallucinating' – in other words, misleading users by presenting inaccurate outputs as factual and reliable information (IBM 2024). Law firms and legal tech companies try to combat this phenomenon by fine-tuning pre-trained generative AI systems – that is, training the model on a controlled set of legal data, including legislation and case-law. The fine-tuning of the model improves its performance – for example, by enhancing the accuracy of its outputs and their relevance for legal professionals, thus mitigating the risk of hallucinations that can compromise the effective management of a case before courts. More recently, retrieval augmented generation (RAG) allows for more accurate and updated responses by generative AI systems, resulting from the retrieval of relevant information sourced from external and current knowledge (e.g. the narrow legal database of the firm or the company), rather than static knowledge encoded during the fine-tuning process (IBM 2023).

The increasing use of generative AI systems by lawyers is raising several questions relating to the ethical duties of competence and confidentiality. On the one hand, lawyers need to stay up to date on knowledge or skills that render them more efficient and effective in the representation of their clients, including digital literacy skills. On the other hand, lawyers risk disclosing their clients' confidential data to third parties – for example, by inserting confidential data when prompting generative AI systems or by giving them access to files containing confidential data. The disclosure of confidential information without the clients' consent compromises the duty of confidentiality, since third parties, including the systems' providers, are now in possession of the data and can process them according to their own purposes.

To understand the impact of generative AI systems on these two ethical duties, the paper seeks to answer the following question:

'What is the impact of the use of generative AI systems by lawyers on the ethical duties of competence and confidentiality?'

The methodology in this paper involves the examination of the professional conduct rules of competence and confidentiality under the Internal Regulation of the Luxembourg Bar Association,⁴ the members of which were questioned through a survey and through interviews.

First, a survey was addressed to lawyers working in Luxembourg to understand their experience using generative AI systems to automate or augment their daily tasks. The survey focused on the duties of competence and confidentiality and how those are impacted by the use of generative AI systems. An anonymous online questionnaire, containing both closed and open-ended questions, was distributed to members of the Bar Association of Luxembourg (Barreau de Luxembourg) through Microsoft Forms, in English and in French. The twenty-eight responses to the

²Recital 97, European Parliament and the Council of the EU (2024) Artificial Intelligence Act.

³Recital 99, European Parliament and the Council of the EU (2024) Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). Available at <https://op.europa.eu/en/publication-detail/-/publication/dc8116a1-3fe6-11ef-865a-01aa75ed71a1>.

⁴Conseil de l'Ordre du Barreau de Luxembourg (2013) Règlement Interieur de l'Orde des Avocats du Barreau de Luxembourg. Available at www.ccbe.eu/fileadmin/speciality_distribution/public/documents/National_Regulations/DEON_National_CoC/FR_Luxembourg_RIO.pdf.

questionnaire were analysed in part automatically, through the generated visualisations (graphs and charts) by Microsoft Forms for responses to closed questions, and in part manually through thematic analysis of the data, having been coded per set of (open) questions and according to assigned labels.

Second, four interviews were addressed to representatives of two law firms active in Luxembourg and two legal tech companies active in France and Belgium, all of which have designed and/or developed fine-tuned generative AI systems for legal professionals. The aim of the interviews was to understand the perception of the designers and/or developers of these systems on the regulation and evolution of the legal profession. The interviews were semi-structured, and their results were processed also through thematic analysis.

The paper concludes with remarks on whether the professional conduct rules on competence and confidentiality need to adapt due to the increasing use of generative AI systems by lawyers. It is claimed that client-centric approaches should ultimately be promoted by Bar Associations, aiming at raising the quality of legal services. Furthermore, the confidentiality of clients' information must be protected through the diligent outsourcing of third parties, the informed consent of clients in the processing of their data, and the training of lawyers on the safe use of generative AI tools.

2 Maintaining competence and confidentiality in the era of generative AI

In Luxembourg, lawyers are organised under the 1991 Law on the Profession of the Lawyer,⁵ which foresees the self-regulation of the profession by the 'Conseil d'Ordre' of the Bar Association of Luxembourg (Article 19). Therefore, lawyers are bound by the Internal Regulation of the Bar Association of Luxembourg,⁶ but they are also considered to be under a contractual relationship with their clients, the breach of which may lead to remedies and damage claims (The Bar of Brussels 2013, p. 11). Lawyers in Luxembourg are further subject to the duty of professional secrecy under Article 458 of the Criminal Code, in principle facing imprisonment or fines for the disclosure of secret information.⁷

2.1 The duty of competence

Under Article 2.4.4 of the Internal Regulation, a lawyer must not take on a case when he or she lacks the necessary skills (e.g. language competence) or diligence to manage it. It is not further explained what 'necessary' would entail, but it can be assumed that this depends on the circumstances and the complexity of the client's case. To acquire the necessary skills, lawyers must have the appropriate level of professional education and life-long learning opportunities, including awareness of the latest technological advancements (Council of Bars and Law Societies of Europe 2019, p. 10). Lacking the necessary information technology (IT) expertise, lawyers could reach out to IT experts to help them strengthen their knowledge and skills in the use of digital systems.

Lawyers should arguably acquire the necessary digital literacy to effectively use generative AI systems that can assist them in managing a case more quickly, consistently and with higher quality of outputs. At the same time, lawyers should be aware of the limitations of these systems, including hallucinations, and avoid relying on their outputs when doubting their validity, accuracy or

⁵Ministère de la Justice (2024) Loi du 10 août 1991 sur la profession d'avocat (Version consolidée applicable au 01/02/2024). Available at <https://legilux.public.lu/eli/etat/leg/loi/1991/08/10/n3/consolide/20240201>.

⁶Conseil de l'Ordre du Barreau de Luxembourg (2013) Règlement Interieur de l'Orde des Avocats du Barreau de Luxembourg. Available at https://www.barreau.lu/wp-content/uploads/2023/04/RIO_Version-consolidee-au-23.08.2022-002.pdf.

⁷Ministère de la Justice (2024) Code pénal (Version consolidée applicable au 08/03/2024). Available at https://legilux.public.lu/eli/etat/leg/code/penal/20240308#art_458.

relevance. This is particularly important for the legal sector, since lawyers cannot competently represent their clients if they rely on defective and unreliable generative AI systems and if they do not verify their outputs. These systems should not be viewed as replacing lawyers because they cannot reproduce the normative judgments involved in legal reasoning or assume accountability for their outputs (Wendel 2019, pp. 28–29).

2.2 The duty of confidentiality

Confidentiality of clients' information generally denotes that lawyers must keep secret and not disclose to third parties any information provided to them by their clients, including correspondence and advice (The Bar of Brussels 2013, p. 1). The protection of confidential information fosters a candid legal communication between the lawyer and the client and helps the lawyer access information that is important for the representation of the client.

Similarly, the duty of professional secrecy demands that the lawyer keeps secret any information related to the client and his or her case, as regulated in Title 7 of the Internal Regulation. The duty covers all types of information, including any advice, verbal communications, correspondence and financial dealings with the client (Article 7.1.3). It has been argued that the duty of professional secrecy also covers information acquired by the assistants of lawyers before, during and after the proceedings, as long as it concerns the client's legal position (The Bar of Brussels 2013, pp. 11–12). The advent of generative AI chatbots assisting lawyers with their tasks begs the question of whether the responses to users' prompts can be considered confidential, if the prompts themselves concern the client's case. The medium of information might not be relevant, so that even information in electronic format, such as electronic correspondence between lawyers and clients, can be considered confidential.⁸ Confidential information can also be found in less obvious sources, such as the metadata of a document.

The duty is regulated as a matter of public policy, being absolute and unlimited in time; however, it can be limited when, among other bases, it is in the client's best interest and the client allows this limitation (Article 7.1.4). In the absence of the client's awareness of, and consent to, the processing of their data through its insertion in generative AI systems, lawyers may breach the duty of confidentiality by disclosing their clients' information to third parties, notably the service providers. This can be done either by inserting confidential data in the generative AI system through their prompts or by giving access to internal files containing confidential information. For example, a lawyer seeking to draft a template of a contract could enter the details of his or her client into the system to generate a tailored result or could give access to internal files so the system can source the (confidential) information necessary for the drafting of the contract.

The disclosure of confidential data is not only a risk arising from the use of generative AI systems but can be found in every cloud-based service. The lawyer must have the infrastructure necessary to permit him or her to conform with the duties posed by the Internal Regulation (Article 1.2). Yet, cloud-based solutions are preferable since they can reduce the operational costs of a law firm, related to the purchase, management and maintenance of digital systems and their servers (Council of Bars and Law Societies of Europe 2012, p. 3). Furthermore, cloud computing applications guarantee access from multiple devices and multiple places, which can increase engagement and responsiveness by lawyers (Council of Bars and Law Societies of Europe 2012, pp. 3–4). Nevertheless, the use of cloud-based solutions can lead to the disclosure of data to service providers, which can compromise the protection of the confidential and personal nature of clients' data. Moreover, the storage of data in the cloud—rather than on a local device or private server—could expose it to a higher risk of unauthorized access to the service provider's servers (Council of

⁸See, for example, Application no. 74336/01. *Wieser and Bicos Beteiligungen GMBH v. Austria*. Court (Fourth Section). 16 October 2007. Available at <https://hudoc.echr.coe.int/eng?i=001-82711>, para. 48.

Bars and Law Societies of Europe 2012, p. 4). These two risks may be exacerbated when the data centres, which are potentially interlinked in a network of servers, are located outside of the European Union (EU) and the European Economic Area, where different standards of data protection are applied.

Lawyers must ensure that external persons with whom they co-operate or collaborate also comply with the duty of professional secrecy (Article 7.1.5). External persons could be providers of Internet, digital office products, cloud-based services and generative AI solutions. The question then arises on whether external providers offering generative AI systems to lawyers and law firms are under an (indirect) duty to respect the confidentiality of the data of lawyers' clients. In similar situations where the duty of professional secrecy is extended to all associated lawyers of the law firm (Article 7.1.6), the supervising lawyer may include clauses in the employment contract to concretise the extended duty (The Bar of Brussels 2013, pp. 336–37; Council of Bars and Law Societies of Europe 2021, p. 15). Something similar could be implemented in the service contract with the external providers.

3 The adoption of generative AI systems by legal professionals in Luxembourg

The advent of generative AI systems in the domain of legal services can in theory have a profound impact on the duties of competence and confidentiality. To understand how generative AI systems are in practice integrated in law firms and used by individual lawyers, a survey was released to lawyers active in Luxembourg that use or are exposed to these generative AI systems. The results of the survey reveal the impact of generative AI systems on the duties of competence and confidentiality as perceived by the surveyed lawyers.

3.1 The use of generative AI systems in legal practice

The first part of the survey focused on the types of generative AI systems available to lawyers active in Luxembourg, the tasks that they automate or augment and the efficiencies that these systems afford in their legal practice.

3.1.1 Generative AI systems used by lawyers and targeted tasks

Fifty-four per cent of the respondents reported on using OpenAI's ChatGPT.⁹ After ChatGPT, the most popular response (25 per cent) was that lawyers were not using any generative AI system in their daily work. Only one respondent stated that he or she uses an internally developed (by his or her law firm) AI tool that is based on ChatGPT.

The primary task with which the largest set of respondents (39 per cent) get assistance is the drafting of various types of documents. Most answers referred to the drafting of generic, non-legal documents, notably e-mail messages.¹⁰ However, certain respondents admitted to the assistance of generative AI systems for the drafting of legal output, including (contract) clauses and framework agreements.¹¹ It is questionable if public LLMs, such as ChatGPT, can perform well in drafting legal documents that have a specific structure and use of terminology. Therefore, it is imperative that users carefully draft their prompts to attach more context to their queries, so that the LLM can

⁹Apart from ChatGPT, popular public generative AI systems and other types of AI systems were Microsoft's Copilot and DeepL. Also mentioned, although receiving one answer each, were the (generative) AI systems of Butterfly, Predictice, Reverso and LexisNexis.

¹⁰Other generic documents mentioned were memos, appraisal forms and marketing communications.

¹¹Apart from drafting, several respondents (25 per cent) get assisted with the refinement of the language used in documents, presumably to render them more grammatically and syntactically correct but perhaps also to confer to them a more concise and legalistic tone.

provide more accurate and relevant responses (prompt engineering), that must in any case be verified by the lawyers.

The next most common task reportedly augmented by generative AI systems (36 per cent) is legal research. By ‘legal research’, respondents denoted the search for legal information, such as court decisions, and more generic information, such as news and foreign practices (in the legal field). Legal research is also taken to mean the extraction of information from a specific document, the explanation of (legal) terms, and brainstorming (on legal matters). In any case, generative AI systems are not like search engines, as one respondent characterised them. Search engines return a list of websites containing the desired information based on a ranking system, while generative AI systems directly compose a response based on the next probable word, not always attaching the necessary references. Consequently, the risk of hallucinations is equally present in this context, although it can be combatted through the verification of the outputs and the application of methods to increase the accuracy and relevance of responses, including fine-tuning and RAG.

Another common task augmented by generative AI systems (32 per cent) is the translation of documents. This result was also apparent from certain software programs mentioned in the preceding question, such as DeepL and Reverso. It is not surprising that translation is one of the most automated tasks by generative AI systems, since lawyers in Luxembourg work in a multilingual environment.¹² Still, given that the legal language contains complex terms with nuances and differences according to the jurisdictional context, public LLMs, such as ChatGPT, may not provide the most accurate and reliable translation, rendering the verification of the output once again essential.

3.1.2 The added efficiency of generative AI systems in legal practice

Most respondents have been using generative AI systems for less or about one year, which is to be expected given that the most popular LLM (ChatGPT) was launched in November 2022. Sixty-four per cent of the respondents report that they are more efficient in their work while using generative AI systems to augment their daily tasks – that is, the drafting and translation of documents and the conduct of legal research, as reported by respondents in the previous question. When asked to elaborate, many respondents echoed the following statement:

‘It helps me work faster. Also, I think that I can use my expertise more efficiently when working on a first AI-generated draft.’

Respondents of the same opinion added that they have to only focus on the verification of the output and apply the necessary corrections to produce a satisfactory document.

The 36 per cent minority stated that they do not perceive any efficiency in the available generative AI systems for legal professionals, either because they are not using them at all or because of the reported instances of hallucinations that require constant verification of the generated text. One respondent characteristically said that:

‘Most answers are very general and require tuning. They also, rarely, are plain wrong (hallucination problems) so the efficiency of creating a block of text is severely counterbalanced by the need to review what has been produced.’

It is interesting that while the majority views verification and fine-tuning of the responses as just another step to render an output reliable, the minority deems them as limitations of generative AI systems that reduce their usefulness for legal professionals. Perhaps the latter category of

¹²The official languages of Luxembourg are Luxembourgish, German and French, while lawyers might also communicate with clients from all over the world.

respondents would like these systems to offer stronger automation capabilities in order to use them in addition or in lieu of available office applications for the facilitation of their daily tasks. Nevertheless, the risk of hallucinations and other deficiencies do not (yet) allow for the full automation of tasks and necessitate human oversight over the system.

Only 25 per cent of the respondents receive training on the use of generative AI systems. When asked to elaborate, some of them reported that internal training by their law firms is provided, while one respondent pointed to online tutorials and books as sources of training. As to the content of the training, respondents mentioned being trained in prompt engineering, compliance issues and best use cases. Apart from those that are not using generative AI systems at all, the majority of respondents may not be interested in training opportunities due to the user-friendly interface of most systems. The system often takes the form of a chatbot, as is the case with ChatGPT, adopting a human-like and conversational tone in its exchanges with users that get quickly acquainted with the model. These anthropomorphic features must not, however, deceive users in relying solely on the model's generative capabilities, rendering imperative the training of lawyers on the best uses of the generative AI systems.

The same two reasons may justify why only 57 per cent of respondents strive to stay informed on updates related to the generative AI systems. Respondents consult mainly media sources, including press releases and published articles. Other respondents stay informed through conferences, internal communications, social media (LinkedIn) and their network.

3.2 Maintaining competence in the era of generative AI

The second part of the survey asked respondents how the increasing adoption of generative AI systems impacts their duty to competently represent their clients.

3.2.1 Duty to remain competent by using generative AI systems

In the question on whether lawyers must use generative AI systems to remain competent when representing their clients, respondents were divided in half. The 50 per cent of lawyers responding positively mainly pointed to time and cost efficiencies, since Generative AI systems can reportedly automate or augment 'routine' tasks so that lawyers can spend more time in tasks that require analytical and communication skills. Other lawyers highlighted the competition among law firms, stating that lawyers must adapt to the new working practices of the legal profession involving the use of digital technologies, in order to not stay behind developments in the legal field.

The 50 per cent of lawyers responding negatively claimed, for the most part, that generative AI systems do not have an added value for legal services. More specifically, they claim that AI systems are not capable of human and legal reasoning to understand the nuances of a case and apply the law to the facts. A respondent characteristically stated that:

'The profession of a lawyer is to offer solutions tailored to a client's needs. This involves considering not only purely factual elements but also emotional and sociological factors. Some of these elements cannot be taken into account by an AI [system]' (translated from French).

This argument may seem to contradict the viewpoint of the other 50 per cent, as mentioned directly above, but both essentially highlight the same point, i.e. that generative AI systems can facilitate only routine tasks, while tasks requiring analysis, reasoning, or communication remain the domain of lawyers.

Another common response was that AI systems often provide inaccurate responses and thus need constant verification efforts, which renders them unreliable as an assistive tool and their use irrelevant for the competence of lawyers. One respondent further explained that:

‘[Wh]en you use generative AI systems you don’t learn at all or at least a[s] much as if you do the research or the writing for yourself, what is especially relevant for young lawyers.’

This line of reasoning corresponds to the responses in the question of whether generative AI systems offer efficiencies that add value to legal services. Many respondents believe that generative AI systems have limited (generative) capabilities and necessitate constant verification efforts, characteristics that reportedly reduce their usefulness for legal practice and minimise their impact on the competence of lawyers.

3.2.2 *Pressure to use generative AI systems and the duty to inform clients*

Eighty-six per cent of the respondents claimed that they do not feel any pressure to use generative AI systems to automate or augment their daily tasks. Only 14 per cent answered yes to this question, stating as sources of pressure the competition in the legal industry and the need for time and cost-efficiency. It is notable that half of the respondents of the survey marked those elements as reasons why lawyers *must* use generative AI systems to remain competent in their legal practice. A possible explanation for this inconsistency might be that respondents took the question to mean whether someone, for example their employer, pressures them to use generative AI systems and thus did not recognise market competition and demand for more efficiency as sources of pressure.

Sixty-one per cent claimed that lawyers do not have a responsibility to inform their clients when they use generative AI systems for case management, sharing the view that:

‘[W]e use tons of different tools in our practice to complete our knowledge and help us solve cases. The AI is just an additional tool which we are responsible as professional[s] to trust, use or not. Lawyers who only want to work with books and fax can do so without telling it to their clients . . .’

At the same time, some respondents belonging to this majority were more modest by mentioning that if confidential and/or personal data of clients are processed through generative AI systems, then a duty to inform clients may arise. One respondent suggested that clients should know how lawyers spend their working time, as using generative AI systems for case management might reduce billable hours. These answers coincide with the minority of 39 per cent believing that there is a need for clients to be informed on the use of generative AI systems by lawyers in order to uphold the values of transparency of legal services and confidentiality of clients’ information. One respondent explained the value of transparency by stating:

‘[T]he contract concluded between a client and a lawyer is a contract concluded *intuitu personae*, meaning it is based on the personal (and professional) qualities of the contracting party . . . If lawyers rely on AI, the profession loses value, and each lawyer becomes interchangeable, resulting in an evident decline in work quality.’ [translated from French]

The reason for this coincidence of views between part of the majority and the minority might be that these lawyers are more digitally literate to understand the parameters of the processing of data inserted through the prompts and the consequences on the duty of confidentiality. Inserting any type of data associated with a client in the generative AI system raises the chances of disclosure to third parties, such as the system’s provider, notwithstanding the fact that similar disclosures may occur when lawyers use other (cloud-based) digital systems in their daily practice. In addition, this group of respondents proved to be more aware of the duty to obtain the client’s consent, according to the duty of professional secrecy under Article 7.1.4 of the Internal Regulation.

3.3 Confidentiality and the processing of clients' data

The third part of the survey concerned the perceived impact of generative AI systems on lawyers' duty of confidentiality. Respondents were questioned on how they are processing their clients' data, if at all, through generative AI systems and whether they apply any compliance and/or security measures.

3.3.1 The processing of clients' confidential data through generative AI systems

Seventy-nine per cent of respondents believe that the processing of clients' confidential data through generative AI systems compromises the duty of confidentiality. The most common justification was the lack of data protection guarantees, including confidential and personal data. Respondents highlighted the risk of breaches when data are collected, stored and further processed by the system's provider for the improvement of its performance. This risk is reportedly exacerbated when the provider is located outside the EU where fewer data protection guarantees might apply, as is the case with public LLMs like ChatGPT.

Other respondents claimed that there could be a compromise of confidentiality given the lack of control over the data once inserted into the generative AI systems. Respondents used phrases such as 'I don't know how they work' or 'it's not clear how data are used' to denote the lack of transparency on the functioning of generative AI systems. This is likely due to the strict protection of proprietary rights over AI systems and concerns about reputational risks from potential breaches, for which system providers may be held accountable. Another reason could be the inability to make generative AI systems explainable due to the complexity of their inner workings that make them difficult to convey in a way that is easily understood by humans. Law firms may also be responsible for the lack of transparency since they do not provide training and information on the best use of generative AI systems to their lawyers, as respondents confirmed earlier.

Twenty-one per cent of respondents stated that clients' confidentiality is not compromised. Some respondents claimed that the generative AI systems they use are integrated into their law firm's secure environment, rather than being freely accessible on the web. Therefore, internal security measures are implemented, including the system's approval by the firm's Data Protection Officer, the storing of data in private servers rather than the cloud and due diligence procedures to ensure that the system's provider offers guarantees to protect confidential data. Other respondents claimed that the compromise of the confidentiality of clients' data depends on how lawyers use the system – for example, on whether they include their clients' data in their prompts or whether they are using public instead of fine-tuned generative AI systems. Consequently, these respondents are in agreement with the 79 per cent majority, concerning the potential risk of data protection breaches when confidential data are processed through generative AI systems. One respondent notably said that as long as no personal data of the client is inserted into the system, the confidentiality of the client's information remains intact. This is a misconception, given that any information concerning the clients' case is confidential, according to Article 7.1.3 of the Internal Regulation, including non-personal data. Lawyers might associate the processing of data through digital systems with *personal* data protection rules but not extend these concerns to the protection of all types of confidential data.

3.3.2 *Compliance and security concerns surrounding confidential data*

The concern of respondents for the compromise of their clients' confidentiality when using generative AI systems is reflected in the 93 per cent of answers stating that lawyers do not insert their clients' data in the systems.¹³ Several respondents justified their unwillingness to process clients' data through the system because of the lack of transparency as to how these will be processed. The 7 per cent minority claimed that they insert confidential data in the system either because there are protection guarantees by the law firm or they simply accept the risk when they upload documents that might include confidential data.

The majority, 79 per cent, stated that they do not follow any legal or regulatory requirements in relation to the processing of clients' confidential data through generative AI systems. Only 21 per cent replied that they apply one of the following rules: data protection legislation, the duty of professional secrecy, the AI Act or internal policies [of their law firm]. This is inconsistent with the earlier majority of 79 per cent of respondents concerned about breaches of the confidentiality of clients' data when these are processed through generative AI systems, a concern which presupposes knowledge of the duty of professional secrecy that bounds lawyers under Title 7 of the Internal Regulation. Lawyers may believe that some of the above compliance requirements burden their law firms and the systems' providers, when in most cases lawyers themselves are data controllers.¹⁴ This means that lawyers are responsible for compliance with data protection rules because they have the choice between using generative AI systems or more conventional (digital) means to complete their tasks and can decide on the purposes of the processing of their clients' data.

3.4 *The future of the legal profession*

The final part of the survey concerned the foreseen evolution of the legal profession due to the increasing adoption of generative AI systems by lawyers and the changes that respondents expect in the regulation and practice of their profession.

3.4.1 *Improvements on generative AI systems towards efficiency*

Most respondents (36 per cent) were reluctant to propose any potential improvements of the generative AI systems to help them increase their efficiency at work. In any case, 32 per cent of lawyers envisioned improvements concerning the reliability of the systems' responses, so the outputs can be more accurate, updated, and accompanied by references. Moreover, 29 per cent of respondents desired new and improved features in generative AI systems, including the recording of timesheets, responses to complex legal questions and more automation in the drafting of (standard legal) documents. A common theme throughout these responses was the need to maintain human oversight over the performance of generative AI systems by delegating only repetitive tasks and not tasks that require legal reasoning. One respondent stressed this point by asking for more awareness of the implications of generative AI systems:

'I would like to have a better education of the lawyers and the people in general about the importance to be sceptical with generative AI systems and our personal data in general.'

¹³It has to be reminded that approximately 25 per cent of the total number of respondents do not use generative AI systems at all.

¹⁴See, for example, Articles 4 (7) and 24 (1), European Parliament and the Council of the EU (2016) Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). Available at <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

Only 11 per cent of the respondents asked for better compliance of these systems with confidentiality and data protection rules, for example by restricting access of these systems only to a pre-defined set of documents.

In the question of whether generative AI systems can replace lawyers or certain functions of the legal profession, 64 per cent responded negatively. The majority of these respondents stated that the emotional and social intelligence of humans cannot be imitated by AI systems, especially concerning the lawyer-client relationship. For example:

‘[H]uman interaction is important to feel and help the client; his questions might be close ended, the answer must always encompass a more holistic or systemic approach to the question[.]’

Other respondents referred to the principles of the professionalism of lawyers that make clients trust them and rely on their advice, and to the principle of human oversight over the functioning of digital systems and their outputs. Some respondents finally pointed to the nature of the legal discipline, namely the ‘nuances of applicable regulation, interpretat[i]on of the law and applicability to each case’ that reportedly cannot be imitated by algorithmic systems but only by humans that have been trained to perform legal reasoning.

The 36 per cent of respondents that answered positively to the question stated that generative AI systems can replace lawyers in the completion of repetitive tasks, including information retrieval and the drafting of standard documents. This is not necessarily contradictory to the opinion of the majority, since it is claimed that only repetitive tasks can be automated, while it might be implied that tasks that require analytical and communication skills are not viewed as automatable. Other respondents pointed to the potential replacement of some categories of legal professionals, including assistants, paralegals and young lawyers by generative AI assistants.

3.4.2 *The adaption of professional conduct rules*

The final questions in the survey concerned the possibility of adapting the professional conduct rules on the duties of competence and confidentiality to respond to the impact of generative AI systems, as described by respondents in the previous sections. Concerning a possible modification of the duty of competence, respondents were divided in half. The first half answered positively, claiming that there must be clear rules on the permitted uses of generative AI systems, notwithstanding the need for the training of lawyers on their best uses. The other half answered negatively, stating that the current rules are sufficient and that ultimately lawyers are responsible for verifying the systems’ outputs and applying them appropriately in their legal practice.

Concerning a possible modification of the duty of confidentiality, 54 per cent of respondents claimed that there is no need for such an alteration because lawyers remain responsible for the protection of the confidentiality of their clients’ data, so that they must simply avoid inserting them in generative AI systems. On the contrary, 46 per cent of respondents claimed that a modification is needed to make clear which uses of generative AI systems are permitted and to facilitate the outsourcing of the development of AI solutions for legal practice, provided that due diligence on the external service providers is implemented. It is notable that in both questions respondents are divided on whether current professional conduct rules need to be adapted or not. This polarisation shows the impact that generative AI systems currently have on the legal community and may be a sign of the evolution of the legal profession in the (imminent) future.

4 The development of generative AI systems by law firms and legal tech companies

The survey results provided an overview of the perception of lawyers as *users* of generative AI systems in their legal practice. To complement these findings, an additional series of interviews on

the *designers and developers* of these systems was realised. Besides the different roles between the two targeted groups, the examined generative AI systems are also different. While the majority of the surveyed lawyers reported on using public LLMs, mainly ChatGPT, the interviewees develop fine-tuned generative AI systems that are trained on a controlled set of legal data. The interviewees' insights demonstrate the process of the gradual integration of generative AI systems in law firms and the anticipated impact on legal ethics and the legal profession.

4.1 Overview of examined generative AI systems

All the examined generative AI systems are based on generative pre-trained transformer (GPT) models. Two respondents stated that their systems are either based on ChatGPT or have a similar user experience to it. All generative AI systems are produced by legal tech companies, although the IT departments of the interviewed law firms have an active role in their development and oversight. The sources of the systems' training data can vary. Two respondents claimed that they train their systems on public information available on the Internet, either as open data or as restricted access data available on subscription-based websites hosting legal content. One legal tech company manages to acquire training (legal) data by public authorities following negotiations. In general, respondents avoid using confidential data to train the examined generative AI systems, which are updated frequently (in most cases weekly) with current information and improved functionalities.

The tasks commonly performed by the examined generative AI systems are the translation, drafting and refinement of documents, as well as legal research.¹⁵ More specifically, lawyers may use these systems to draft content, ranging from e-mail messages to contract clauses, and refine their language likewise. Respondents reported on the efficiencies of their developed systems. Law firms highlighted improvements in the speed and quality of drafting contract clauses, refining written text or summarising documents. These improvements extend beyond the English language to other foreign languages for lawyers using the systems, such as French and German. Legal tech companies respectively underlined the facilitation of complex contract management through easier access to knowledge and the improved reliability of the systems' responses through the attachment of references.

Nevertheless, the examined generative AI systems have their own limitations, according to respondents. The most reported limitation is hallucinations, which can especially affect legal research when lawyers ask open (legal) questions and expect to receive accurate and updated answers. Respondents called for more awareness of the risk of hallucinations, so that lawyers always verify the systems' responses and adapt their content according to the context in question.

A different limitation reported by a legal tech company is the lack of context specificity, in particular:

'We do not have the technology yet to . . . write a simple prompt and then there is a document that comes out of it . . . [W]e deal with complex contracts where . . . you need a lot of context to make the output very relevant . . . [W]e have found that not to be the case or [that] the user experience is not there just by writing a simple prompt.'

Generative AI systems are often criticised for not being able to understand the context of a prompt, which then leads to generic responses that are not useful for users. This may be particularly the case for the legal domain, where public LLMs, such as ChatGPT, are not trained specifically on legal data and cannot therefore understand the background of the prompt in order to respond to legal questions. The company attempts to combat this issue through evaluation

¹⁵Other tasks that are reportedly automated or augmented include the summarisation of documents and the transcription of events, such as meetings.

scenarios, where legal experts prepare specific prompts which developers can then insert into the system to test different system setups. The responses to the prompts are then compared to the legal experts' model answers to understand how well the generative AI system can perform in the legal context and if there is need for further fine-tuning on legal data.

Respondents raise awareness of the limitations of generative AI systems by providing training to lawyers in the form of demonstrations and case scenarios on the best uses of the systems. These practices also aim to combat any reservation of lawyers on the use of generative AI systems. As one law firm claimed, lawyers might fear that they are being surveilled by their colleagues or third parties while using these systems and that they might be subsequently criticised for asking 'stupid' questions or violating applicable regulations. One law firm also gives lawyers the opportunity to submit feedback as users of the system, so the firm can identify the need for any improvements. Another law firm instructs its lawyers on how to prompt the generative AI system to receive relevant responses to their prompts – for example, by refining the language input in the prompt to tweak the system. One legal tech company attaches references in the form of hyperlinks at the end of the chatbot's answers to encourage the verification of the system's results.

4.2 Maintaining competence in the age of generative AI

Law firms stated that lawyers must use generative AI systems in their daily work due to their reported efficiencies in automating or augmenting certain tasks. One law firm emphasised the professional responsibility of lawyers by stating that:

'With or without technology, you will always have good lawyers and bad lawyers . . . [I]f you are a good lawyer, you don't take for granted something that is being spit out from a machine. You double-check and do your work and you think about what is being . . . written, put your name on your memos, so that what you are really giving to the client is your work product.'

Respondents representing legal tech companies generally agree with the law firms on the need for lawyers to use AI tools to remain competent in their work. They added that there is no risk of replacement of lawyers due to the need for critical thinking to assess the systems' outputs. Both companies focused on the state of competition in the legal services market and claimed that generative AI systems can make lawyers more efficient timewise, so they have additional time to spend on more strategic tasks. One company stated that there are efficiency gains for both smaller and larger law firms in terms of competition. Smaller law firms can fine-tune generative AI systems to fit their field of expertise, while larger law firms can expand the applicability of these systems to fit their multiple areas of expertise.

Interviewees were divided on whether lawyers should inform their clients about whether and when they use generative AI systems. One of the law firms replied that if lawyers are using the system for tasks such as the summarisation or translation of legal documents, and they do not process client information, then clients do not need to be informed. The other law firm stated that clients need to know how lawyers are spending their working hours, so the latter can be charged accordingly for their services. However, it stated that there is no need for a disclaimer on every activity undertaken by lawyers, given the trust inherent in the lawyer-client relationship. It is notable that they both failed to mention the need for the consent of clients when their confidential information is disclosed to third parties (Article 7.1.4 of the Internal Regulation), as is likely the case with data inserted in the generative AI systems and later shared with the systems' providers. This omission may be attributed to the limited processing of clients' data through the examined systems, as seen in the following section.

4.3 Protecting the confidentiality of information when inserted in generative AI systems

Law firms stated that they do not in principle process their clients' data through generative AI systems, including for training purposes. However, one law firm claimed that exceptions are made based on the consent of clients or on legitimate interests when consent cannot be relied upon because it can be removed at any time. This is an interesting statement, given that lawyers need to obtain the consent of their clients every time they disclose their confidential information, including through generative AI systems (Article 7.1.4 of the Internal Regulation). Moreover, one legal tech company admitted to accessing lawyers' document management systems, following the extension of ethical walls and permissions to the company by law firms. Access to these files permits the company to improve their systems' understanding of users' queries and their performance in locating and retrieving the desired contract clause in lawyers' databases. Snippets of these contracts may be stored on the company's servers, but they do not usually contain identifiable data. In any case, the company underlined that lawyers have the possibility to restrict access to their documents to avoid the processing of confidential data.

Respondents also refrain from disclosing confidential data to third parties. Both legal tech companies assured of the non-disclosure of the data of lawyers' clients to third parties. However, one law firm admitted its frustration in not being able to use digital tools without the possible sharing of data with third parties:

'Today you can't find any law firm . . . that doesn't use a solution which is cloud-based . . . [F]or the moment we are in the grey zone because to a certain extent we have an obligation not to share that [clients'] information but at the same time we need to do our job and you can't do that job without any technology.'

This response highlights the intrinsic limitations of cloud-based solutions regarding data protection and the desire for clearer rules on the obligations of law firms, as analysed in the final section of the interviews' analysis. The other law firm claimed that the only information they share with third parties concerns the user experience, including users' feedback. Exceptionally, they might provide third parties with only-once, supervised access to data in a human language format to acquire advanced technical support. In any case, they avoid disclosing any input or output data of the generative AI system, including users' prompts. In addition, their (private) servers are based in Europe, so they are in full control of the data and there are no data transfers to third states outside the EU. This is also the case with one of the legal tech companies.

Respondents comply with several regulations, depending on the jurisdiction(s) where the law firms and legal tech companies offer their services. Since all respondents are active in the EU, they comply with the General Data Protection Regulation, which regulates, among other issues, the disclosure of personal data and the implementation of impact assessments. Other regulatory regimes include International Organization for Standardization standards and the Information Security rules. Respondents also abide by the professional conduct rules of the Bar Association of Luxembourg. One law firm underlined that lawyers are ultimately responsible for the (use of the) outputs, even if (outsourced) providers of generative AI systems are also responsible when handling confidential information for the development and improvement of their systems.

All respondents apply security and organisational measures to restrict the processing of confidential information, especially their accessibility and reuse by third parties. One common measure among all respondents is the encryption of (confidential) data. Both legal tech companies delete lawyers' data, either when they cease to be customers or when the data constitutes personal data. Data silos are also applied by a legal tech company, presumably to enable the isolation of confidential data in a particular part of the company's systems that is secure from unauthorised access. Both law firms reported that new vendors are thoroughly reviewed through information security and risk assessment processes before they are contracted. One law firm internally restricts

access to data on a need-to-know basis, while no single individual controls the whole process. The firm also applies strict auditing on each step of the system's performance, so there is compliance by design according to the current legal and technological framework.

4.4 The future use of generative AI systems in law firms and the relevance of legal ethics

Respondents revealed their desire to improve some features of their generative AI systems, a process that necessitates greater access to lawyers' data. In general, generative AI systems for legal professionals can improve their performance overtime depending on the number and types of data inserted in the model. If the task is to analyse a contract, then the more contracts are processed by the system, the more it will understand the legal language and structure of the contracts, and it will be able to accurately and consistently analyse them.

One law firm already implements a contract metrics tool that refines the language of contracts, having been trained on lawyers' databases that host such documents. A legal tech company seconds this need to access lawyers' databases, so the generative AI system can check the documents and identify any patterns revealing illegal activities in legal matters such as mergers and acquisitions. Another legal tech company aims to develop an anonymisation tool that masks text excerpts, without missing the context that is necessary for the searchability of the document. The company is also interested in adding features that help capture collective knowledge and expertise within the law firm, so every member of the law firm benefits from this knowledge, even if individual members responsible for its collection are no longer part of the firm. One law firm foresaw the adoption of voice recognition features, which can help to automatically refine a prompt through a dynamic verbal interaction with the user.

As to the future of the legal profession influenced by the integration of generative AI systems in law firms, all respondents agreed that lawyers are not going to be replaced by generative AI systems. One law firm claimed that these systems cannot replace lawyers because they cannot replicate the human relationship (developed between the lawyer and the client) and have not yet overcome certain limitations, such as hallucinations. The respondent instead foresaw the arrival of the "augmented lawyer" that uses generative AI systems to improve his or her performance in daily tasks, including the drafting of texts. This is reportedly an important skill that lawyers must possess in order to compete with digital (AI) tools that support self-representation through access to legal information. Another represented law firm found necessary the adjustment lawyers' training to focus on how to delegate their repetitive tasks to generative AI systems, presumably to speed up the process of case management.

Respondents were not convinced that professional conduct rules, notably competence and confidentiality, should be modified due to the increasing use of generative AI systems in legal practice. Both legal tech companies, instead, underlined the need for innovation within law firms so that lawyers are assisted by these systems, notwithstanding the constant verification of their outputs. One law firm stated that legal ethics should not adapt, since they are in principle basic, straightforward and easy to understand. This opinion reflects what is currently the practice in Bar Associations around the world – that is, the publication of guidelines instead of amendments or new rules on the use of generative AI systems by lawyers.

The other law firm expressed its desire for new rules to address how to deal with the sharing of data with third parties. In particular:

'[E]ach law firm has to determine what they can put in place to say that they've done everything they could to actually keep all the information that they receive from their clients confidential. We're not like in the banking sector, where they have like proper provisions and where it's actually clearly scripted with whom a professional of the banking sector can share information . . . For the moment, for lawyers, we don't have let's say legal framework . . .'

In Luxembourg, the Financial Sector Law provides exemptions to the duty of professional secrecy, including for the outsourcing of activities to an external service provider (Mbayi and Bulach, Unknown). The consequence of this exemption is the free transfer of clients' data to external providers – that is, with no obligation of professional secrecy vis-à-vis the provider. This happens either when there is a service agreement with the external provider or when the client has accepted, under the terms and conditions, the outsourcing of activities for a particular set of data.

On the contrary, under Articles 7.1.5 and 7.1.6 of the Internal Regulation, external collaborators and members of the law firm are not exempted but must comply with the duty of professional secrecy. If an exemption similar to the financial sector was to be extended to the legal sector, then lawyers could freely share (confidential) data with outsourced providers of generative AI systems, with neither party being at the risk of breaching the duty of professional secrecy. This freedom would likely allow law firms to maximize the efficiency of their legal practice by fully leveraging the benefits of the generative AI systems they use. Still, the compromise of the confidentiality of clients' data would be probable and additional measures should be agreed between the two parties to safeguard these data.

5 Client-centric legal services in the age of generative AI

Lawyers and developers alike have shown through their responses a keen interest in using generative AI tools to be more efficient in their everyday work, while also being aware of confidentiality issues arising from the potential disclosure of their clients' information to third parties through these systems. Given this risk to clients' confidential information, the focus of regulators should be shifted from raising the efficiency and competitiveness of legal practice to raising the quality of legal services through the use of generative AI systems.

Client-centricity aims at high-quality, affordable, variegated and innovative legal services in order to enhance access to justice for (prospective) clients, whose interests must be in the center of the regulator's mind (Semple 2015, p. 243). Self-regulatory regimes established by Bar Associations need to look beyond lawyers' interests, while preserving the principles of professionalism and independence that underly the legal profession. Regarding the augmentation of legal practice by generative AI systems, Bar Associations need to be proactive in setting down rules on the best use of these systems by lawyers to address clients' interest in keeping their data confidential. These rules could take the form of either comments attached to the existing articles or separate guidelines distributed to all members of the Bar Association.

5.1 Informed choices by clients

Article 7.1.4 of the Internal Regulation of the Luxembourg Bar Association states that professional secrecy can be limited when it is in the client's best interest *and* he or she allows it. Therefore, the client must consent to his or her data being processed through generative AI systems. The Internal Regulation further states that, prior to consenting, clients need to be informed by their lawyers on the nature of the shared information and on its recipients, most likely the systems' providers. Depending on the given case, there might be further information to be communicated so that the client can make an informed decision.

In the absence of specific rules in the Internal Regulation, the General Data Protection Regulation provides further conditions for a valid consent when the personal data of clients situated in the EU are being processed.¹⁶ For the consent to be valid, it needs to be freely given, specific, informed and unambiguous (Article 29 Working Party 2018, pp. 5–18).¹⁷ The client needs to be informed on the elements that are essential in providing his or her consent. First,

¹⁶Articles 3 (1 and 2) and 4 (1), European Parliament and Council of the EU (2016) General Data Protection Regulation.

¹⁷Article 7 (4), European Parliament and Council of the EU (2016) General Data Protection Regulation.

clients need to know how the use of generative AI systems by lawyers contributes to the management of their case, since lawyers can potentially spend less *billable* hours working on a particular case and produce the same or better results due to the insights provided by these systems. They must also be informed when they are expected to pay higher legal fees because their lawyer must in turn pay an ‘out-of-pocket’ fee to use a particular generative AI tool for case management,¹⁸ especially when their clients instruct them to do so (ABA 2024, pp. 11–14).¹⁹ Second, clients need to understand what constitutes confidentiality, which types of data or information are confidential and for which purposes they are processed. Third, clients need to know which (third) parties have access to their confidential data and whether these may be transferred to third states – for example, when the providers’ servers are located outside of the EU. Fourth, clients must know the consequences of potential breaches of confidentiality and the applicable security and organisational measures for the proactive and reactive management of such breaches.

In addition, clients need to have control over whether their data can be processed through an generative AI system, without facing any negative consequences on the quality of the legal services offered for not providing their consent. Consent is not considered freely given if the provision of legal services is tied to the processing of data that is not necessary for the performance of the contract. It could be argued that the processing of clients’ information through generative AI systems is not necessary, considering the possibility to use of less risky digital systems for the completion of legal tasks. The consent needs to be specific, so that the client consents to the processing of his or her data through the system for a specific purpose and for this purpose only – for example, for the drafting of a contract. Furthermore, the lawyer should share any information necessary for the provision of consent in a clear and plain manner, avoiding long texts with legal or IT terms, so that the client provides his or her consent unambiguously, in the form of a written statement or a clear affirmative act.

If clients finally decide to not provide their consent for the disclosure of their data to the providers of generative AI systems, lawyers cannot prompt the system with confidential information or share access to files containing such data. This obstacle was identified by a law firm asking for an exemption to the duty of professional secrecy, so that lawyers can share any data with outsourced companies, including the systems’ providers, and increase the efficiency of their work without breaching their professional duties. On the one hand, this could be against a client-centric approach by prioritising the interests of lawyers in improving their performance rather than protecting clients’ confidential information. On the other hand, the constant technological advancements and the increasing integration of generative AI systems in law firms may in time normalise the sharing of clients’ data through them, as is the case with other cloud-based office applications. In the end, clients may also benefit from more time and cost-efficient legal services due to the use of generative AI systems that might provide important insights into their case and contribute to its effective representation before courts. In any case, a possible exemption from the duty of professional secrecy does not undermine the responsibility of lawyers and their law firms to implement security measures for compliant data processing procedures, as analysed in the next section.

5.2 A safe and compliant generative AI system design

Following the client’s valid consent for the disclosure of confidential information to the providers of generative AI systems, lawyers can insert this data into their prompts or provide systems with

¹⁸For example, for the generation of insights that can contribute to a successful litigation.

¹⁹The bill can also be incurred for hours spent on the training of the lawyer relating to the use of the tool; clients should not, however, be charged additional hours when the use of the LLM is part of the routine of a lawyer’s practice (e.g. when it is used to check the grammar of a document, such as any other cloud-based office tool).

access to files containing confidential data. Lawyers must then supervise the providers of these systems to ensure that they respect the confidentiality of the clients' data (Article 7.1.5 of the Internal Regulation). Alternatively, an exemption from the duty of professional secrecy for both the lawyer and the system's provider would allow for the free disclosure of any type of data, including confidential information, notwithstanding the need to apply security and compliance measures for the processing of data.

In either case, law firms must specify the parameters of the processing of clients' information through the generative AI system in the service contract (Council of Bars and Law Societies of Europe 2012, pp. 5–9). More specifically, when law firms perform due diligence on the system's provider, they need to review the experience of the provider in terms of years and certificates or licenses, its reputation regarding client satisfaction and inexistence of security incidents and the location of the company and its servers whether in or outside the EU. During the negotiation of the service contract, law firms and lawyers must review the data protection policies and other security measures implemented by the provider. There could even be explicit clauses clearly and unambiguously defining what constitutes confidential data of the lawyers' clients and stating the limits of their processing, prohibiting in any case their processing for the training of the generative AI system.

Measures to further prevent the repurposing of data²⁰ for the training of generative AI systems or for other related purposes can include the identification of confidential data within existing and future documents, including in the documents' metadata, and the restriction of access to these documents by third parties or other staff of the law firm. Any confidential data that are deemed irrelevant for the client's representation or that relate to concluded cases can be deleted. Some law firms could opt for anonymisation or pseudonymisation of data instead of deletion. However, neural networks underlying these systems can often enable reidentification – for example, through the linkage of two or more records related to the clients' cases (Article 29 Data Protection Working Party 2014, pp. 11–12).

Monitoring and reporting obligations should be established, including impact assessments focusing on the likelihood and severity of data management and security breaches. These breaches can include cyber-attacks targeting the models' training data, their algorithm or the prompts themselves. The model's training data can also be biased so the system returns biased outputs to the user, or the prompts might be 'injected' by malicious actors, so that the system follows other instructions other than those defined by the user (Rehberger 2023). There is also the likelihood of unauthorised access to the confidential data during their processing by generative AI systems either by third parties or even by people working in the same law firm, who may be prohibited from accessing this information due to an ethical wall or who might (intentionally or not) use this information in the representation of another client (ABA 2024, p. 7). The impact assessments must be made on a case-by-case basis, taking into consideration the technical parameters of the tool, the task(s) that it performed and the client's case in the context of which it was used. In the event of a breach of confidentiality, law firms must have in place an effective mitigation plan and notify clients of the breach of their confidential data and the applicable mitigating measures.

Designers should strive to involve legal experts in this procedure and inform them of any limitations that might implicate their daily work (Kapoor, Henderson, and Narayanan 2024, pp. 4–7). Developers must in turn attempt to make generative AI systems explainable to empower users to make informed choices, rather than overrelying on the system due to its anthropomorphic characteristics (Carli and Calvaresi 2023, pp. 5–15), for example when the system's interface takes the form of a chatbot. The AI Act establishes further obligations to the providers of generative AI

²⁰Article 5, paragraph 1 b and c, European Parliament and Council of the EU (2016) General Data Protection Regulation; Longer periods of storage and avoidance of deletion can also risk the accuracy of the data since they remain unchanged in the dataset and might not be updated regularly or at all (Article 5, paragraph 1 d).

systems.²¹ If these constitute a general-purpose AI system, as public LLMs are, then the provider needs to implement an updated, technical documentation of the model, including its training and testing process, and make this and other related information available to other providers wanting to integrate the model into their AI systems (Article 53, para. 1). They should also draw up a policy for the protection of copyright and other related rights, while making available a detailed summary of the content used for the training of the model. A fine-tuned generative AI system may still be characterised as a general-purpose, if a general-purpose AI system has been integrated into it (Recital 100). For example, a law firm or an outsourced company could integrate OpenAI's GPT models into their own model to serve a variety of purposes, such as generating contracts, summaries and translations.

For high-risk AI systems, providers need to follow an additional set of obligations, including the implementation of a design that allows for human oversight and ensures an appropriate level of accuracy (Articles 8-15). It is arguable if generative AI systems for legal services can be considered high-risk. To the extent that the processing of clients' data through these systems can interfere with their rights to privacy and personal data protection – for example, when their personal data are used for the generation of a contract without their consent, then the system could be characterised as high-risk.²² Generative AI systems designed for more administrative tasks, such as summaries of law, would probably not fall under this category.

5.3 Trustworthy and innovative legal services

For the establishment of trust between clients and their lawyers, the latter need to be liable for actions that have repercussions on their clients' rights and interests, including misapplications of generative AI systems. Delegating tasks to a machine without verifying the accuracy and reliability of its outputs, at least to a reasonable extent, is a sign of incompetent representation, since lawyers are not demonstrating a sufficient understanding of the limitations of these systems and/or do not perform the necessary steps to mitigate any potential risks, such as hallucinations.

Stricter disciplinary penalties might be a measure not only to sanction breaches of professional conduct rules but to also set an example to other members of the Bar Association regarding the best uses of generative AI systems. However, it is a measure of last resort, when more preventative measures should be preferable. The most obvious would be the training of lawyers on the best uses of generative AI tools, as a life-long learning measure to sharpen the digital literacy skills of lawyers and ultimately maintain their competence. As seen in the results of the survey and the interviews, there is usually no formal training opportunities provided by law firms in Luxembourg. To increase the consistent application of best practices, training events could be offered by Bar Associations themselves in different formats and combining different modules.

More specifically, training courses could be offered in person, online or in a hybrid mode. The courses could be one-time events or consist of multiple sessions, where attendants can receive certification by the completion of the course. Lawyers could be trained in both the theoretical and practical aspects of using generative AI systems in their legal practice. Theoretical aspects could raise awareness of the efficiencies and limitations of these tools and their impact on professional conduct rules. One such session could tackle in a simplified way how generative AI systems function, mainly in terms of how they are trained and how they generate a response to users' prompts. Emphasis should be given on how the system processes (personal and/or confidential) data and on the likelihood, severity and impact of data breaches compromising the confidentiality

²¹The providers of LLMs in this case can be either outsourced companies developing the model and/or placing it in the market or law firms themselves which are putting a model in service under their name; Article 3 (3), European Parliament and the Council of the EU (2024) Artificial Intelligence Act.

²²An AI system is classified as high-risk when it has 'a significant harmful impact on the health, safety and fundamental rights of persons in the Union'; Recital 46, European Parliament and the Council of the EU (2024) Artificial Intelligence Act.

of clients' data. Lawyers can then be more vigilant when they allow generative AI systems to access their clients' data or when they insert them in prompts, especially when the client has not consented to such processing.

Practical aspects that could be covered in workshops (perhaps given in collaboration with IT experts) could entail the implementation of best practices in different case scenarios, depending on the legal task at hand. A useful practical session could focus on prompt engineering, a method to effectively formulate requests to the generative AI system to receive more accurate, reliable and relevant results. Lawyers can benefit from learning how to craft clear and concise prompts, incorporating relevant context through related texts or examples, and breaking down complex queries into distinct parts to allow the system to process them more effectively ([OpenAI Platform, Unknown](#); [Microsoft Learn, Unknown](#)). Above all, lawyers should master the discipline to always verify and fine-tune the outputs of generative AI systems according to the context in question, as well as apply the indicated measures to mitigate a materialised risk, such as a data breach, including the notification of the designated parties in the law firm.

6 Conclusion

Generative AI systems offer new ways to assist lawyers in their daily working practice by augmenting a range of tasks, from the translation of a document to the drafting of a contract. Lawyers must keep abreast of technological developments that contribute to their work in representing their clients. At the same time, the use of generative AI systems that can process clients' confidential data for unauthorised purposes and share them with third parties, such as the systems' providers, can compromise the duty of professional secrecy.

According to the results of the survey and the interviews, both users and developers of generative AI systems are focusing on augmenting the drafting of various types of documents, including contracts. Users are mostly resorting to public LLMs, mainly ChatGPT, while developers are releasing their own fine-tuned models. Probably not all lawyers have access to fine-tuned models through their law firms, since it is expensive to develop from scratch such a system or even outsource it. Public LLMs are much more affordable and easily accessible on the web but lack context-specificity and accuracy due to their training on generic sources from the web.

Both users and developers were positive that the introduction of generative AI systems in the workplace will not replace lawyers, rather augment their performance while managing their clients' cases. However, users of the systems were still concerned that some functions and categories of legal professionals could be replaced. Moreover, both groups of respondents were aware of the risk of confidentiality breaches whenever they allow generative AI systems to access and otherwise process their clients' information. As a result, they do not, in their majority, process clients' data for the purposes of training the systems or receiving answers to their prompts.

The respondents of the survey and the interviews were also divided on the same issues. First, a significant number of users and of developers was reluctant to proactively inform clients as to the use of generative AI systems in their law firm, stating that this should be the case only when clients' data are being processed through the systems. It should be once again reminded that lawyers have an obligation to obtain their clients' (valid) consent under Article 7.1.4 of the Internal Regulation. Second, not all users and developers supported the modification of the rules on competence and confidentiality. Some stated that the current rules suffice, while others asked for new rules clarifying the permitted uses of generative AI systems.

One important point of controversy between the respondents of the survey and the interviews regards the need for lawyers to use generative AI systems to remain competent in their legal practice. While all developers of these systems were positive that lawyers must use them due to their efficiencies, only half of the users were of the same view. The other half of users stated that current generative AI systems present instances of hallucinations and require constant verifications that do not render them useful for daily tasks. One possible justification for this

difference of views is that developers are probably more willing to promote their products among lawyers by highlighting their efficiencies.

Another point of controversy concerns the training of lawyers on the use of generative AI systems. Most users stated that they do not receive any training on the use of these systems by their law firms, while developers mentioned that they offer demonstrations on their best uses. It is likely that the interviewed law firms and legal tech companies feel bound to offer training since they want to promote their own systems, while law firms that simply accept the use of public LLMs by their members do not recognise such an obligation on their part. A final point relates to the potential improvements of generative AI systems currently offered in the market. The majority of users asked for generative AI systems that better comply with the duty of confidentiality – for example, by limiting access to lawyers’ files containing confidential data. On the contrary, developers asked for the possibility to use more (confidential) data to be able to offer systems that are more fine-tuned and tailored to users’ needs.

The way forward would involve a client-centric approach, where the use of generative AI systems by lawyers would depend on the interest of clients in high quality legal services. Bar Associations must adopt rules, in the form of comments or guidelines, that target three aspects of client-centricity. First, the rules should require the valid consent of clients for the processing of their confidential data through generative AI systems. Second, the rules should demand the inclusion of specific clauses on the protection of clients’ confidential data in the contracts signed between lawyers and providers of generative AI systems. Third, the rules must promote life-long learning opportunities on the use of generative AI systems, so lawyers can remain competent while representing their clients and respectful of the confidentiality of their clients’ information.

Acknowledgments. The present paper has been written in the context of the author’s doctoral research, funded by the Luxembourg National Research Fund (Fonds National de la Recherche) under the PRIDE programme (PRIDE 19/14268506). The paper was first presented at the International Legal Ethics Conference 2024 (ILEC 10) at Amsterdam Law School (University of Amsterdam). The author is grateful to the Bar Association of Luxembourg (Barreau de Luxembourg) for circulating the survey among its members and to all the members that participated in the survey. The author would also like to thank the representatives of the law firms and the legal tech companies that shared their insights and opinions during the interviews.

Competing interests. None.

References

- American Bar Association (ABA)** (2004) Formal opinion 512 on generative artificial intelligence tools. Available at https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf
- Article 29 Working Party** (2018) Guidelines on consent under Regulation 2016/679. Available at <https://ec.europa.eu/newsroom/article29/redirection/document/51030>
- Article 29 Data Protection Working Party** (2014) Opinion 05/2014 on anonymisation techniques (0829/14/EN WP216). Available at https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf
- Carli R and Calvaresi D** (2023) Reinterpreting vulnerability to tackle deception in principles-based XAI for human-computer interaction. In Calvaresi D *et al.* (eds), *Explainable and Transparent AI and Multi-Agent Systems. EXTRAAMAS. Lecture Notes in Computer Science*. London: Springer, 14127, pp. 249–269.
- Council of Bars and Law Societies of Europe** (2012) CCBE guidelines on the use of cloud computing services by lawyers. Available at https://www.ccbe.eu/fileadmin/speciality_distribution/public/documents/IT_LAW/ITL_Position_papers/EN_ITL_20120907_CCBE_guidelines_on_the_use_of_cloud_computing_services_by_lawyers.pdf
- Council of Bars and Law Societies of Europe** (2019) Charter of core principles of the European Legal profession. Available at https://www.ccbe.eu/fileadmin/speciality_distribution/public/documents/DEONTOLOGY/DEON_CoC/EN_DEON_CoC.pdf
- Council of Bars and Law Societies of Europe** (2021) Model code of conduct for European Lawyers. Available at https://www.ccbe.eu/fileadmin/speciality_distribution/public/documents/DEONTOLOGY/DEON_CoC/EN_DEONTO_2021_Model_Code.pdf
- Hu K** (2023) ChatGPT sets record for fastest-growing user base – analyst note. Available at <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

- IBM** (2023) What are large language models (LLMs)? *IBM*. Available at <https://www.ibm.com/think/topics/large-language-models>
- IBM** (2023) What is retrieval-augmented generation? *IBM*, Available at <https://research.ibm.com/blog/retrieval-augmented-generation-RAG>
- IBM** (2024) What are AI hallucinations? *IBM*. Available at <https://www.ibm.com/topics/ai-hallucinations>
- IBM** (2024) What is cloud computing? *IBM*. Available at <https://www.ibm.com/think/topics/cloud-computing>
- Kapoor S, Henderson P and Narayanan A** (2024) Promises and pitfalls of artificial intelligence for legal applications. *Journal of Cross-disciplinary Research in Computational Law* **2**, 1–12.
- Mbayi M and Bulach M** (Unknown) Outsourcing versus Luxembourg professional secrecy. International Bar Association. Available at <https://www.ibanet.org/article/BE078227-1EAB-40C8-88A9-BCA19583E476>
- Merken S** (2023) New York Lawyers sanctioned for using fake ChatGPT cases in legal brief. *Reuters*. Available at <https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>
- Microsoft Learn** (Unknown) Introduction to prompt engineering with GitHub Copilot. Available at <https://learn.microsoft.com/en-us/training/modules/introduction-prompt-engineering-with-github-copilot/>
- OpenAI Platform** (Unknown) Prompt engineering. Available at <https://platform.openai.com/docs/guides/prompt-engineering>
- Rehberger J** (2023) Real-world exploits and mitigations in LLM applications (37c3). Available at <https://www.youtube.com/watch?v=qyTSOSDEC5M>
- Sample N** (2015) *Legal Services Regulation at the Crossroads: Justitia's Legions*. Cheltenham: Edward Elgar Publishing Limited
- Spataro J** (2023) Introducing Microsoft 365 Copilot – your Copilot for work. *Blog Microsoft*. Available at <https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>
- The Bar of Brussels** (2013) *Professional Secrecy of Lawyers in Europe*. Cambridge: Cambridge University Press.
- Wendel WB** (2019) The promise and limitations of artificial intelligence in the practice of law. *Oklahoma Law Review* **72**, 21–49.