AI Revolution in Law

On November 30, 2022, ChatGPT was launched. The public was stunned by the capabilities of ChatGPT and all of a sudden anyone who could type could access the seemingly near-human power of artificial intelligence. How was the arrival of this type of generative AI received in the legal industry?

Pinsent Masons is a international law firm, originally founded in England more than 250 years ago, with over 26 offices across 4 continents. Two years ago the Amsterdam office has been launched. Alastair Morrison, Head of Client Strategy and Board Member of Pinsent Masons LLP and Andre Walter, Head of Data Law Solutions in the Netherlands, reflect on the arrival of ChatGPT at Pinsent Masons.

Alastair: We first talked with clients about this when ChatGPT2 came out. That was around January of 2020. There was quite a big gap between ChatGPT2 and ChatGPT3. ChatGPT3.5 came out, and suddenly, version 4 came out.

When ChatGPT came out on the market, the impact was astonishing. I personally was astonished at how it had moved from ChatGPT2 to ChatGPT3 and I can remember the first kind of questions I asked ChatGPT and what came out of it. I would ask ChatGPT to tell jokes, to write a poem, to explain a complicated legal structure in the style of a particular comedian. But when you start seeing the kind of scope and capability of it, we are all amazed by what it all can do. The acceleration and the knowledge

of Generative AI and examples such as ChatGPT has been growing exponentially since.

As a result, when Generative AI arrived, we at Pinsent Masons started very quickly thinking about what this will do for the service economy and how it's arrival is going to change it.

So, the moment Gen AI / ChatGPT came out Pinsent Masons was already planning on using it?

Alastair: We've been delving into the potential applications, with less focus on specific use cases and more on the broader question of whether this will significantly impact professional services, spanning across various dimensions. For the first time, we're witnessing a technology that effectively harnesses its summarization capabilities and techniques. And this is merely considering the capabilities of ChatGPT3 and ChatGPT4. The true intrigue lies in observing Gen AI's capabilities when these skills are cultivated within a secure environment. That's when things truly become captivating.

However, when you consider Gen AI's overall potential and its possible outcomes, you inevitably start asking fundamental questions: Who will it replace, and who will it complement? Which skills will be enhanced through augmentation, and which will be refined by its assistance? Moreover, which skills are susceptible to automation? Is the use of generative AI in the legal field inevitable?

Alastair: Yes, 100%. Especially if you look at what it's capable of doing and the different use cases of it. For example, performing a legal or contractual analysis.

However, there are many lawyers who'll just freak out about the latter and say: I'm the owner of the information that is used by the AI program. Or they'll say: I can do the legal or contractual analysis myself. And, of course, that's right. But it depends on how good these lawyers are, what kind of skills and what kind of capabilities they've got. But what you have got here - signifying a point where augmentation becomes exceedingly captivating—is a consistent base case and foundational capacity. And of course, there is the hallucination issue that I will come to later. As for example ChatGPT is picking up the sum of everything that has been written on the internet and because a lot of rubbish has been written on it, it is possible that it picks up rubbish. The lawyer however will be able to tell you if it's rubbish very quickly.

The extent to which the generative AI might gather irrelevant information depends on the content it's being trained on and the datasets at its disposal. Obviously, if you're training with rubbish, you're going to get rubbish. But if you are filtering that out and looking at contractual analysis based on particular contracts then there's no room for the rubbish.

The summarization capability of generative AI gets you to the starting point so much quicker than you could have done before. It's like a library of all human knowledge that can give you the information in warp time. And, of course, you might still read something in a book, or you might not agree with the outcome of the summarization as you look at the matter in a different way, but still, it can give you the base case very quickly, leaving more time to truly reflect and advise on the implications and consequences of a particular issue.

So as for now you see it as a good tool to do all the preliminary work which can be replaced by ChatGPT. This brings us to the other side of the story. In the legal field, ChatGPT is often considered as something that one day could possibly replace the legal team. Additionally, according to a recent report from Goldman Sachs 44% of legal tasks could be performed by AI. What do you think of this? Is this fear for full automation among legal experts justified? How does Pinsent Masons approach the balance between integrating AI and preserving human jobs within the legal field?

Alastair: Stepping back, one might question: what kind of change will this bring about? There's so much to sort out in what I often describe as the universe of law which like

our universe is constantly expanding. Major corporations do struggle to get their heads around all the different challenges that they face. And some of the top lawyers in the major organizations or corporations therefore often talk about the Perma crisis or the Poly crisis. So it's a good idea to have a closer look at what some of our teams that we work with are facing at the moment to see how Gen AI can help them. In house lawyers have multiple issues to deal with such as their regulatory environments, compliance, ESG, Digital Transformation, Privacy, Data, Sanctions, Climate Change , and that's all on top of business as usual litigations , corporate deals , employment issues.

As a consequence Chief Legal Officers have forged significantly stronger ties with their organization's CSuite, so a closely intertwined relationship has emerged between Chief Legal Officers, the CTO, and the Chief Information Security Officer (CISO). You're starting to see very complicated issues on legal-risk advisory work, legal-risk management, and risk and compliance, and these issues need to be solved. So, you need a bigger machine that is able to do that. And lawyers, traditionally, have been working with a pick-and-shovel play to try and get these things sorted out. Now, you have this giant excavator that you can use as a tool to sort this stuff out at scale. So for law firms to really properly embrace this, they can answer questions far quicker and at scale.

You already talked about generative AI being capable of performing legal and contractual analysis. Is Pinsent Masons already using generative AI for other day-to-day operations, and if so, how?

Alastair: We must make a distinction between classic AI - machine learning techniques being able to extract big sums of data from documentation and then incorporate them into contracts - and generative AI. Our business and other professional services organizations have been doing classic AI for quite a while already. For example, legal technologists who work with lawyers to produce the required results, particularly in the financial services sector where you've got a vast range of contracting work, including contract remediation work.

One can think of the situation where there's been a legislative change and all the contracts that have been written around that particular space have to be revisited. You can't possibly have individuals going through each of these contracts in order to change them. So, we've

already been using techniques in order to extract, summarize and put into context the right documentation.

We're actually quite cautious around ChatGPT. We told people to have a look at this new technology and to discover it. However, you cannot put everything on ChatGPT. There are rules and regulations around us, trying to work around the relative information of clients. By putting that kind of information into your ChatGPT prompts as a lawyer, you'd breach your privacy obligations as a lawyer. If you put anything around client information onto ChatGPT, it's into the public domain and that way it's gone. Private Gen AI tools is the way the professions are going.

What's the solution for the problem of not using sensitive information?

Alastair: It all revolves around the type of generative AI you employ and the safeguards you implement within a privacy-controlled environment for AI utilization. But we don't want to just sit there doing nothing around all this. You do want professionals within organizations to be exploring and learning what this new technology is all about and asking generic questions and seeing what its capabilities are, but not using it to give legal advice in relation to client related work.

As models evolve through, there are going to be some major tech providers on it and open-source organizations will also integrate with the major players to offer evolving enterprise AI solutions. That's the kind of thing we're looking at now. But now you just want to encourage people to open their minds to what type of potential this has so they become more familiar with it.

André: You see a similar problem with translation programs. You can't just put client information or personal data into a translation program as it may give raise to confidentiality and privacy concerns. So, in many cases you must first strip it down to anonymous templates, before putting it into the translation engine. Therefore, Pinsent Masons adopted its own translation engine that runs in our controlled environment. This is probably something that you can also expect in the future for the ChatGPT kind of services.

Is Pinsent Masons using the same ChatGPT the rest of the world is using or do they have a special program developed?

Alastair: What we're doing in our technology team is looking at different use cases and applications of where one might be investing in the longer term. But we haven't really called them on that yet. There's lots of different things going on, and an organization like us has very good relationships with the major tech providers as we've got lots of licenses with some very big companies. We are really trying to explore and experiment regarding generative AI. In conclusion, we do use the technology that's behind the products of OpenAI but not the products itself, because of their security issues.

Another large international law firm has already integrated a legal AI tool called 'Harvey' into its practice, using it for contract analysis, due diligence and litigation preparation. Does Pinsent Masons has such a tool, and if not, is that something Pinsent Masons will be working on in the near future?

Alastair: We are aware of these developments and are closely monitoring them. I think there is quite a lot of early noise around these kinds of tools. Behind the scenes, we are actively engaged in working with these tools. We know of the research and development going on and looking at what to do. We are taking steps to ensure that Pinsent Masons remains competitive considering these technological advancements.

We already talked about the implications of the integration of generative AI. Recently, Steven Schwartz, a personal-injury lawyer at the New York firm Levidow, Levidow & Oberman used a generative AI, ChatGPT, to help him prepare a court filing. It created a motion replete with made-up cases, rulings and quotes. How does this incident relate to the reliability of the integration of generative AI within the legal field?

Alastair: I think this comes back to the essence of professional services organizations. One can't just use the general open AI and try and get an answer to your question. Of course, you can use it as a starting point to work out, for example, a particular legislative regime.

I think the importance of the foregoing is that you can't jump too quickly to the generated answer. It's all about the security and structure of your data sets. If your data is structured and well assembled, you're in a far better position. Great datasets will include external data like data from the legal publishers. So that you're plugging into trusted data and secured data sets.

As AI technology continues to evolve, for example, a new version of ChatGPT will be released soon, how does Pinsent Masons plan to stay ahead of the competition and remain a pioneer in integrating AI advancements into its legal services?

Alastair: First of all, what you have to look at in the context of your business is what do our clients want? And for that we look at it in three ways.

The first thing is the use of generative AI in the advisory space. So, that's kind of what Andre will do, giving advice to clients on the implications of generative AI from an Intellectual Property perspective, or whatever the perspective might be, on their business. So that's the traditional legal work in the context of generative AI, which is advising clients on use cases, data privacy issues that the client may have.

The second thing is the User Experience (UX) . How do we within Pinsent Masons use generative AI to make us faster and better horses if you like. It's investing in new generative AI models that improve the user experience.

Then the third point is the Customer Experience (CX). Our strategy is to look at what the customer is seeking to achieve? What is the desired outcome for their business? We then leverage the generative AI applications to streamline operations, eliminate inefficiencies, and optimize processes. A lot of our contact with clients entails sitting down with them and working out in the context of their business. What can generative AI potentially do for them and what issues may arise out of it?

In a recent article, The Economist predicts that legal services ultimately will become cheaper because of the rise of generative AI. Thereby legal services will be more widely available for small and medium-sized businesses that previously often struggled to afford them. Do you think legal services will indeed become cheaper because of generative AI, or will the client pay what some call a 'technology fee', so that clients don't expect to get generative AI for free?

Alastair: There are two things going on. There's the corporate world and then there's a democratization of legal services. We'll come back to the corporate world in a minute as to what might be corporate related advice and where that's going to go.

In the context of the democratization of legal services, you can see that what generative AI is able to do is to provide things rapidly, at scale for people who have not been able to afford to get legal advice before. People talk about this systematization and commoditization of legal services and keep on saying that it's not quite as bad as a lawyer might think. One does see good platforms and good consumer advice, lots of them coming from the United States. I think these platforms will get exponentially more powerful, more accessible, easier, quicker, faster and cheaper.

If you look at the percentage of the world's population that has access to legal justice, it is tiny. Therefore, these models can probably and should properly introduce something which is more different from the legal services we know of nowadays. It could become something which will be NGO led. In any way, we will see purposeful generative AI models that would help the democratization of legal services. Meanwhile, one could also ask if this means this will bring access to justice also.

If we get into corporate land and look at what clients want to demand on the delivery of legal services,

investing in generative AI is not cheap. Meanwhile, I think the most consistent and most reliable advice will be delivered when you put the lawyer together with the technology. You see the same in medical research. AI is used to do the first run.

I think within professional services organizations, particularly within law firms, you will see a shift from hourly rates to more lump sum fixed prices based on value. Law firms price work on what it costs them in contrast to what the value is to a client. That shift will move over to what the value is for a client and pricing will be based on that value. That will result in better legal advice, because it has to be valuable legal advice. It will also result in cheaper legal advice in many cases, as you take inefficiency out of the system. However, not necessarily cheap just for the sake of being cheap.

André: What we also see in practice regularly, is that we have to ' design to budget'. For example, we may have a client who wants us to review new developments in a larger number of legislations. At the same time, clients do not want to spend thousands of euros per jurisdiction. This often reduces the scope to a high-level review.

The aim of using generative AI would not be delivering the existing product for half of the price, but producing a deliverable that is ten times more detailed. Generative AI supported advice could deliver much more value for money.

In a recent study by Carnegie Mellon University computer scientists found that it is possible to automatically construct adversarial attacks on Large Language Models (LLM's), such as ChatGPT. When specifically chosen sequences of characters are added to a user's query, it will cause the system to obey user commands even if it produces harmful content. Is Pinsent Masons aware of these attacks? Further-more, is cybersecurity something which that will be considered during the further integration of generative AI into Pinsent Mason's day-to-day practice?

Alastair: Absolutely! I think this also comes back to what I call the 'first-mover loser' ('first-mover' is the competi-tive advantage gained by the initial ('first-moving') significant occupant of a market segment, red.). That really is something you don't want to be. You need to be very thoughtful around your security and privacy when first using these new kinds of technology. Security and Privacy becomes even more important in the context of this, which is why you have to be thoughtful about what you put out to market.

André: I fully agree. The security and privacy, especially at the input site of these generative AI models, one must be careful not to breach any intellectual property or privacy rights.

Alastair: If you look at the future of professional services, one will be looking at the way these organizations train their people. In the past, you would learn a lot and then take your expertise and put it into an advisory context. However, the access to knowledge and expertise is now much quicker and that has implications for how we train and develop professionals of the future.

Another noteworthy point is the emergence of new job roles. Legal prompt engineers, for instance, will likely become prominent. Professionals will learn how to train and guide these AI systems to provide accurate solutions and responses. It's conceivable that lawyers might transition into roles like legal hallucinationsurveyors, where their function involves overseeing instances of potential legal hallucinations. On that note, someone once quipped to me that AI will be responsible for deciphering AI-generated hallucinations – the AI will assess AI's hallucinatory outputs.

What Pinsent Masons is already doing now, is looking with our graduate recruitment teams on what to think about this and what the implications are, what we might want and what these jobs of the future will be.

About the interviewee
Alistar Morrison is a Board member of Pinsent Masons with overall responsibility for the firm's client and sector strategy. Alastair's role is to advise the Board on the objectives, strategies and investments required to sustain the firm's relevance to clients and build long term sustainability, resilience and relevance in the global legal market. Prior to his Board role, Alastair has been the principal construction law adviser on numerous major projects on-shore and off-shore nationally and internationally.