



# Taxing Matters

## Exploring the World of AI with Olivia Dhein

**Alexis**

Hello, and welcome to Taxing Matters, your one stop audio shop for all things tax brought to you by RPC.

My name is Alexis Armitage and I will be your guide as we explore the sometimes hostile and ever-changing landscape that is the world of tax law and tax disputes.

Taxing Matters brings you a roadmap to guide you and your business through this labyrinth.

In case any of you miss any crucial information or just want some bedtime reading, there is a full transcript of this and indeed every episode of Taxing Matters on our website at [www.rpc.co.uk/taxingmatters](http://www.rpc.co.uk/taxingmatters).

I am absolutely delighted to be joined today by Olivia Dhein. Olivia is a Knowledge Lawyer at RPC. Prior to this Olivia was a Professional Support Lawyer at Lexis Nexis and, before that, a Commercial Disputes Lawyer at Freshfields.

I think it is also fair to say that, over the last year, Olivia has become something of an expert in a topic everyone is talking about - at the moment - and that is: 'artificial intelligence', or 'AI'. Olivia has become an expert in the interplay between AI and law and has a thorough understanding of the technological underpinnings and legal ramifications of the technology.

Olivia and I are going to be discussing: "what is AI?", "what is a 'large language model' and how does it work?", "what can AI do and what can AI not do at the moment?", "how is AI likely to affect tax advisors and accountants?" and "is AI here to stay and, if so, what should we all be doing about it?".

Let's get started. Olivia, welcome and thank you for joining me today!

**Olivia**

Thank you and thank you so much for having me, I'm really excited, because this is my first podcast ever!

**Alexis**

Well welcome, here we are! Olivia, what is AI?

**Olivia**

A really good question. You will not be short of news stories on AI as this amorphous, big thing, that everybody is talking about. I find it quite helpful to break it down into talking about 'large language models' as that is what people are usually talking about, rather than the wider field of 'artificial intelligence'.

Chat GPT hit the scene a year and a bit ago - in the public eye - and it is probably the most famous example of a large language model. So, really everybody knows what a 'large language model' is, there're other models out there. So, for example, "LLaMA", which is by Meta, "PaLM" and "Gemini" which are Google/DeepMind, or "Claude" by Anthropic.

But, probably for ease in the conversation, I'll refer to ChatGPT because everybody knows that one.

**Alexis**

What is a "large language model" then and how does it work?

**Olivia**

This is my favourite question actually, because, once you get that explanation, a lot of things just fall into place. A "large language model" - or, I'll call it LLM - is, basically, a highly sophisticated computer model which has learned to understand the structure of human language.

Let's just take ChatGPT. The people who invented ChatGPT took this model and the first task they gave it to do was to read the entire internet, which it did in about a month, which I always find -

---

<b>Alexis</b>	Wow!
<b>Olivia</b>	<p>- blows my mind!</p> <p>I was chatting to someone who is an expert in this area, and I said, "ChatGPT read the internet in a month!", - this was "unbelievable" - and he just laughed and laughed, and he said, "oh, it's just a matter of processing power, if you have the right processing power it can do it in minutes!" So, -</p>
<b>Alexis</b>	Wow!
<b>Olivia</b>	<p>- this just gives you an idea of the breadth and the scale of these models - so, that's why I explained that - but the model went around reading language - to simplify a little bit - it basically mapped out language. It also mapped out images and sound and computer code - but just for simplicity I'll stick with language - and it's the most relevant bit for lawyers anyway. In a way, it went around and turned all the words that it came across - so, imagine every Reddit thread, every blog, every Amazon review, or whatever it could find - mapped out words and how they link together. Basically, turning words into numbers - because it's a computer model. It does that in a way so that concepts that are similar are closer together in its map and concepts that are further apart in meaning are also further apart in this map.</p> <p>It also looks at how words relate to each other in a sentence. So, for example, the word "interest" can have different meanings. The model compared these kinds of sentences, it always looks at all the different relationships between the words and can, basically, work out which kind of "interest" is being talked about. It's a really simplified explanation, but it has created this enormous map of language.</p> <p>Once it had finished that it was trained to answer questions and put together the words. It's effectively a huge prediction machine that runs a very, very complex probability calculations, trying to predict which word comes after which, based on a question that somebody puts into the model.</p> <p>That was the next stage, actually, in how ChatGPT was trained. After it had learned the internet and mapped it all out a team of humans would sit down and evaluate how well the answers matched the questions which is a type of reinforcement learning - so that was an extra layer on top. Then, they did another step, which was to add on so called 'guard rails' - which was just another layer of code, effectively, to stop the model from doing anything harmful. What they were trying to prevent was, for example, a model giving instructions on how to commit a crime, or advocating violence, or anything like that. For me, the basic steps. Fundamentally, that is ChatGPT, or any other 'large language model'.</p>
<b>Alexis</b>	Great! thank you. And so, what can AI do and what can it not do, at the moment - because, obviously, it's all moving so fast, isn't it?
<b>Olivia</b>	<p>It's moving incredibly fast, yes!</p> <p>I was just saying to somebody, the other day, "I have my notes from my talk and every two weeks I cross something out and I just have to change it because it's just moved on!"</p> <p>At the moment mapping the language and predicting words and putting them all together, is giving these models natural language processing, NLP. That is to say, it can draft really well, it can understand text, it can read text, it can summarise text. It can, to some extent, answer questions about a text - a little bit like a child that has gone to primary school and has learned how to read. It has not necessarily learned how to reason and, although - question mark - whether it can do that in the future, maths is not something its particularly good at, but again that is changing.</p> <p>In terms of more basic language tasks, it does a pretty good job. So, for example, if you ask it to summarise text, that is something it can do very nicely! Or draft in a certain style - it can do that very nicely as well!</p> <p>I think another thing that is really important to remember is that this is something that has actually "finished school", it's finished its training. A lot of people get a bit confused about this side of it. So, these models can run by themselves, and you can get information out of them, but they are doing this purely on what they have</p>

---

been trained on and they are extremely complicated calculations of probabilities - which word comes after which - so that's how they spit out the answer – and, sometimes, the answer is not correct. Sometimes it has nothing to do with reality and that is what we call a “Hallucination”. The reason for that is that simply, sometimes, it just calculates probability that happens to not be true in real life. The model isn't malfunctioning, it's just that it has never been asked to check anything against real life, or against the database, or anything like that! It's just doing a prediction of what it thinks the most likely answer to the question will be.

So, that is one risk with these models, “Hallucinations” - they might get it wrong. At the same time, you can hook-up these models to search engines, or a database and then you can ask it to cite a webpage, or cite a document, so you can check where it gets its information from, if you wanted to. So, there are ways around this.

Another thing, that I think is quite important to mention is that, as I said: “it's a prediction machine”. It's a huge simplification, but I find it quite easy to remember that and everything the model does is prediction but it's not a copy of anything else, anywhere else. These models are so unbelievable, in a way, because they, actually, create something new every time you ask them a question, or ask them to do something. They are not taking a copy of something that they have stored in a massive database and giving you the most likely - or most relevant - answer to your question, they are, actually, deciding to give you the exact words it is going to use.

**Alexis** So, if I asked a question on the system and you asked the same question on the system, we would get different answers? Is that what you mean?

**Olivia** Yes. Very likely that we would get different answers.

Let's just say we're using the same ChatGPT account. At first you type the same one in and then I type the same one in, it might think. “oh, it didn't like the first answer to my question so I better offer up the one the next, most likely, probability because that might suit the human better as to what they were looking for”. Again, it's not malfunctioning, it's just an inherent feature of how these things are set up, especially quite useful to know for when you want to use them.

**Alexis** Definitely! So, is it the case then that we are all, effectively, training the system how to respond? Whoever is using the system is training it to respond in a certain way?

**Olivia** Yeah, I mean, it depends which one you're using. If you are using a public version of ChatGPT, yes, that is exactly what they do!

They use what the public are typing in to further refine and train their models because they are still doing research and they are still, obviously, developing; but at the same time, if you are using something that is a product that is set up for companies, usually that is not the case. So, to protect data privacy those sorts of LLM products are usually set up in a way that you don't train the model with the data that you put into it.

**Alexis** Thank you. How do you think AI is likely to affect tax lawyers and advisors and accountants then?

**Olivia** The first thing to say on this topic is you need to work out a strategy of how you think you can use the technology. Because it has so many possible uses, I think the best way to think about it is to think about, “what could I do if I had a machine that can read and analyse language for me?”, “what could I do with that?”. Generally speaking, I think the tax lawyers might have to do that kind of research on a party, or something like that. So, anything that involves large volumes of documents would be useful. But unlike using some kind of ‘keyword search’, or search engine, it is more intelligent because the model understands context. Let's just say I need to find anything that has to do with a “garden”. The model would, conceptually, understand that it needs to offer me things that have to do with flowers, with trees, with butterflies, with insects, small animals – whatever - it will understand that these contexts hang together.

So, this isn't something we've ever, had available to us before.

**Alexis** And so, who do you think is at the forefront in this space at the moment, in terms of AI for tax advisers and accountants?

**Olivia** I think there's a lot of tax laws, I'm sure any assistance would be very welcome. Both Thomson, Reuters, Westlaw and Lexis Nexis have announced that they are bringing to the market LLM driven versions of their

platforms. So, this is quite exciting for lawyers because it means that they have taken an LLM and they have plugged it into their databases, they have done some clever 'fine-tuning' around it to minimise 'Hallucinations' and then they will present it to you to make your research easier.

Whatever your tax law question is, rather than having to search for it and thinking about in terms of search terms, you can use natural language to describe what you are looking for and that might just end up saving you a huge amount of time trying to find the right information - even just from a known text book!

**Alexis** *In your mind, do you think AI is here to stay and if so, what should we all be doing about it?*

**Olivia** The short answer is: "absolutely!", the genie is out of the bottle! I don't see this going anywhere. If anything, I see this accelerating and - to be honest - if anything, I think the world of law is a little bit behind where the tech actually is and the kind of products we can access are not anywhere near cutting-edge and it's quite exciting to see what is happening there and I'm really interested to see more about models that can, for example, reason more than GPT4, for example. GPT4 is really-good at drafting things, giving you answers, but in terms of reasoning, logical reasoning - that's what I'm really excited to see in the future - that could obviously help us as lawyers.

The speed of development is super-fast! I read that the chips that this stuff runs on are going to double in processing power this year and I am expecting there to be all sorts of very interesting effects and much bigger models - and possibly much cheaper running costs which might affect pricing, so that we might see even more useful products being offered up to us. I think it's here to stay - That's a really long way to say that!

**Alexis** *A bit of a hard question to answer, but: quite a lot of people - when they think about AI - they are quite worried about job cuts and things like that, do you think that it's likely to change the tax world?*

**Olivia** I think so, I think there's a huge opportunity - especially for law - because we can get rid of a lot of the drudgery. All that perfectionist checking that we're so proud of doing! We might not have to do it so much anymore and, actually, I'm not too sad about that because we can focus on more interesting legal questions.

I also wonder if these tools, that we can expect to come our way, if they make a lot of the work quicker - in that sense - that we might be able to do more cases, or that legal services will become even more accessible so there will actually be more work because more people will be able to use them. I think it's a huge opportunity.

**Alexis** *We've talked previously about this but what is the current Law Society guidance on using AI?*

**Olivia** We have official 'AI guidance' for the use in English courts and tribunals which came out just before Christmas - I think people missed it, because it was just in the pre-Christmas rush - which was interesting to me. It, basically, warns of the risks of AI and tells judges that they cannot put private information into a public chat box and so on and sets out examples of how you shouldn't use it. But, very interestingly, it also says that judges can use AI in their work and they don't actually need to declare it.

It also says that lawyers can use AI in litigation, for example, and they don't necessarily have to declare it either unless the judge decides that they need to specifically answer that question. I think it's another sign that the courts have realised that this is here to stay and something that is evolving and that we just need to get used to.

**Alexis** *There's been a lot of talk recently in the tax world especially about this case called "Harbour v HMRC" that took place in the first-tier tribunal; all about AI and how fake citations were used. The AI system had thrown out hallucinations. So, would you be able to tell us a bit more about this case and what your view is on it?*

**Olivia** I have strong views on this topic!

Not only this case, but views on this topic! Because that is the one topic that has really gone through the legal press - all through 2023 - is international examples of lawyers citing cases that were, actually, fabricated, or "hallucinated" by ChatGPT and so on, and 'Harbour' happens to be a tax-case, but it is the first English case that I am aware of where this has happened. In this case Mrs Harbour, in her appeal, was given the dubious assistance of a friend who worked in a solicitor's office - and probably expected to be given real cases - and,

unfortunately, she was given references to cases to support her argument on 'reasonable excuse for not notifying her tax liability', that were just fictitious! The tribunal did a very careful exercise to look through the fake case citations and compared them to real cases and found, in fact, that it was highly likely that the cases had been generated by a large language model such as ChatGPT.

I do feel for Mrs Harbour because she had no idea and she was not represented. This is, definitely, not a case of somebody doing the wrong thing, but of being unassisted in their legal research. I think in the decision it is mentioned that she was not aware of legal research websites such as 'Bailii'. So, this is not a criticism of her whatsoever.

It was interesting to me because the tribunal was so careful in analysing it. To me, that is that! So, I think we have the right outcome. The tribunal emphasised and endorsed an international case in the US explaining why harm can come from these fake case citations - not least because names and parties and judges are being cited and that can cause reputational harm because what is being said is not true.

I think the really, really, fundamental point is really important. But, does it tell us much about AI? I don't think so. It is a really important point, I feel, that it is reassuring that this was picked up exactly in the way it should have happened. If, in good time, people start using the Westlaws and the Lexis Nexis of the world that are powered by this technology would this happen? I don't think so, because you would, actually, have a tool that is meant for legal research that uses AI and will most likely give the right result and, even if it didn't it probably has built in checks and balances to avoid Hallucinations, because I am sure those companies have worked very hard to make that happen.

So, to me, these cases where lawyers are very, very upset - it's really just bad research - it's bad legal research that shouldn't have been done that way! Certainly, if you are a legal professional, obviously it's different if you are a litigant in person. Really what we should be focusing on is learning how to use this new technology and see how far we can push good legal research and how we can make it safe and what benefit we can, actually, give our clients in doing it in this new way and of course there will be some hiccups when we learn how to do that, but that doesn't mean we shouldn't do it.

<b>Alexis</b>	So, is the take away message from this then that that case is actually almost like a trainee doing the legal research, nobody more senior checking it, and it going straight into the tribunal and there is actually a mistake in there? is that the overall message - that AI - if it's managed properly can be a really good tool but it shouldn't be just left to its own devices?
<b>Olivia</b>	Yeah, I think that's exactly right. It's a bit as if you are using an assistant in some shape or form and don't check what comes out. In this particular case I think the case citations had been provided with summaries of the cases, so you could actually check quite easily by just sticking on the citation and looking for the full text and, in fact, that is, maybe, what courts will do in future - just require full text citations and they could just require confirmation of which legal research website you have got it from - which is another way of dealing with it - that was mentioned in the judgment as well.
<b>Alexis</b>	Quite a lot of people across the industry have got excited about this case and said, "well, this is a reason why we shouldn't use AI and these are all the problems with it!". But you don't share that view?
<b>Olivia</b>	I don't share that view, no. It doesn't really tell us anything other than that large language models Hallucinate and we've always known that. So, you should never work with a large language model if you don't know how that comes about and what your checks are to prevent you from using the wrong output.
<b>Alexis</b>	So, it's all about the checks and balances that you use alongside the AI?
<b>Olivia</b>	There's a huge opportunity there to get really-great legal advice done much more quickly.
<b>Alexis</b>	And surely the benefits to clients will be potentially huge because costs should be, theoretically, less if AI is able to do a huge amount of the legal research and other tasks, especially the document analysis and things like that?
<b>Olivia</b>	And avoiding errors as well! So, research errors.
<b>Alexis</b>	Yeah, absolutely. Another question I've got is that you mentioned "Lexis" and some of these other providers are going to be launching some new AI tools early this year. How available are they going to be to your

---

average taxpayer who may or may not have representation, or is it still a case that actually they are much better off going to tax advisors and lawyers to get that advice because this stuff probably won't be available to them?

**Olivia** I do think that "Lexis" and "Westlaw" solutions are very much for professionals, I don't think they would be terribly useful for lay-people, but also, maybe from a cost perspective. I am not quite sure where the pricing is, but I imagine it's not something that you would just get for £3 a month!

**Alexis** Well exactly! So, it would probably still be inaccessible to your average taxpayer, in which case the real benefit will be to the advisors and the lawyers and the accountants who will be able to use the AI to then benefit the end customer, i.e. the client. It will probably still be like that for a while I would imagine - until it is rolled out widely if it even needs to be. Is there anything else you wanted to say on AI generally?

**Olivia** Yeah, there was one thing: have a play with it!

If you haven't already, if you are a lawyer and you are a bit scared of it, don't be! Try out ChatGPT, or Google Bard, whichever model takes your fancy, because I think it is a really, good exercise - at the moment – to, just, try and get used to using it and see how these models behave. Don't put any personal data in there, don't put anything confidential in there obviously! Just have a play around with it, see how it comes back with answers and testing it. Also, so I tested ChatGPT on "Small-talk" the other day because I wanted to see how good it is at that: how many new topics it can bring it into the conversation and so on, so whatever you fancy. It's just a good illustration and when you see how quickly it comes up with a draft it gives you, it reminds you of that scale of data and computer architecture that sits behind these models, they are enormous. They do spit out some pretty interesting stuff.

**Alexis** So, that's the takeaway everyone, for your homework is to go and mess around with ChatGPT!

Well, thank you so much Olivia, we have certainly learned a great deal about AI today! Unfortunately, that's all we've got time for in this month's episode. Thank you again to Olivia for today's podcast. You can contact Olivia at [olivia.dhein@rpc.co.uk](mailto:olivia.dhein@rpc.co.uk).

A full transcript of this episode together with our references can be found on our website [www.rpc.co.uk/taxingmatters](http://www.rpc.co.uk/taxingmatters).

If you have any questions for me or for Andrew or any topics you would like us to cover in a future episode, please do email us on [taxingmatters@rpc.co.uk](mailto:taxingmatters@rpc.co.uk). We would love to hear from you.

If you like TaxingMatters why not try RPC's other podcast offering, Insurance Covered which looks at the inner workings of the insurance industry hosted by the brilliant Peter Mansfield and available on Apple podcast, Spotify and our website. If you like this episode, please take a moment to rate, review and subscribe and remember to tell a colleague about us.

Thank you all for listening and talk to you again soon

---



RPC is a modern, progressive and commercially focused City law firm. We have 97 partners and over 700 employees based in London, Hong Kong, Singapore and Bristol. We put our clients and our people at the heart of what we do.