

Tech Talks Podcast

The Global AI Regulatory Maze: Impacts on Tech Transactions & Governance

April 18, 2024

Announcer

Welcome to Mayer Brown's Tech Talks Podcast. Each podcast is designed to provide insights on legal issues relating to Technology & IP Transactions, and keep you up to date on the latest trends in growth & innovation, digital transformation, IP & data monetization and operational improvement by drawing on the perspectives of practitioners who have executed technology and IP transactions around the world. You can subscribe to the show on all major podcasting platforms. We hope you enjoy the program.

Julian Dibbell

Hello and welcome to Tech Talks. Our topic today, artificial intelligence. Once again, today we are focusing on the emerging landscape of AI regulation across the world and its implications for technology transactions and internal governance. I'm your host, Julian Dibbell. I'm a senior associate in Mayer Brown's Technology & IP Transactions practice. I'm joined today by Ana Bruder, Oliver Yaros and Arsen Kourinian. They are all partners in Mayer Brown's Cybersecurity & Data Privacy, as well as our Artificial Intelligence practices. Ana is based in Frankfurt, Oliver is in our London office, and Arsen sits in our Los Angeles office. Very happy to have you all here today. We have a lot to get through. So let's get into it.

Ana, I want to start with you because we need to talk first about the EU AI Act, the biggest news right now in AI regulation. Can you tell us roughly what the EU AI Act is about, what the current status is and the next steps in its adoption?

Ana Bruder

Yes, absolutely, Julian. Thank you for having us on your podcast again. So the EU AI Act was adopted on the 13th of March by the European Parliament. And that's it. The EU has done it. We are just a very few steps away from having the very first comprehensive AI law in the world. And to be honest, we're quite proud of that.

So the next step is that the council is going to formally endorse the text and then it will be formally adopted. It will be published in the official diary of the EU and that is expected to happen any time between April, so in the coming weeks, and August this year.

Julian Dibbell

Okay, so what happens then after formal adoption of the EU AI Act? When does it start applying?

Ana Bruder

There will be a staggered implementation of the provisions. So within six months of adoption, the provisions on banned AI systems will start applying. So, for example, social scoring or using biometrics to select people that will not have access to services, for example, that is banned, that will start applying already six months within adoption.

Then one year after adoption, the provisions on general-purpose AI systems will start applying. And really the bulk of the obligations under the EU AI Act will start applying two years after adoption, some three years after adoption. In particular, those that relate to high-risk AI systems. And there's a subcategory of high-risk AI systems that is very intertwined with you – product safety legislation, and those are the ones that will take longer to start applying, so three years.

Julian Dibbell

Okay, three years. That's a nice long horizon, but it sounds like you're saying that some provisions of the EU Act might actually start applying this year. Is that right?

Ana Bruder

That is right. So if anyone is using systems that will be banned, they should be already looking to stop using them because once this applies, there will be high fines, very high fines for noncompliance with banned provisions, banning provisions, I mean. And then there's more than that. There's actually stuff already going on right now to implement the EU AI Act.

For example, the authorities that will supervise and enforce the new rules, they are already being set up right now. Which by the way, there's a word that the EU AI Act uses to refer to that, to the authorities that will supervise and coordinate. The EU AI Act calls that governance, much to the confusion of global audiences that may be more used to the word governance in its US use of the word, which refers rather to the internal policies and procedures that companies implement to manage AI-related risks and also more broadly to the obligations that apply to companies when they're developing or using AI. But if we focus really on the UAI Act use of the term governance that is already happening right now. So the AI office, for example, was already established as a body within the European Commission. And so there were some job postings open until end of March. We're looking for experts to help with that. And right now, the AI office is just awaiting you member states to nominate the national competent authorities that will integrate it. Spoiler alert, data protection authorities are very interested in the job.

Julian Dibbell

Okay, so that's governance in the European sense, turning to the US, United States context and how we use the word governance here, what's the impact there of the EU AI Act? How is that going to affect companies, including US companies?

Ana Bruder

The EU AI Act, just like the GDPR and many other pieces of legislation in the EU, has extraterritorial applicability. That means it will not apply only within the EU territory. So any companies throughout the world that may be using AI, where the output of the AI system is intended to be used in the EU, that will trigger applicability of the UAI Act. So let me try to give you a very concrete example. You're a US company, you're using an AI tool in the context of employment, say to filter CVs, right? To make recommendations for a vacancy based on decisions that were made in the past. If that vacancy is in the EU, that means that the output of the system will be used in the EU. That will be enough to trigger applicability of the UAI Act.

So that means that for each AI system being developed or used, companies will actually need to assess several aspects in order to determine first the applicability of the UAI Act, and then as a second step, which obligations, if any, arise to them. So let me tell you, just go through those questions, the key questions that companies need to be asking themselves is, what is their role with regard to the AI system? Are they a provider, a deployer, an importer or a distributor? Second, is the system general-purpose AI or not? Because the set of obligations depends on that classification. If it is general-purpose AI, is there systemic risk or not? And there are provisions to help you assess that, of course, in the UAI Act. And if it's not a general purpose AI system, what is the level of risk? Also that, so all of these aspects are key for companies to assess which obligations, if any, really apply to that specific AI system. I now would say providers of general-purpose AI and high-risk AI systems, in particular, they will have really the most extensive obligations. And we're talking conformity assessments, very extensive compliance obligations that relate, for example, to cybersecurity to data governance more broadly, including privacy, quality management, there needs to be a risk management system, extensive technical documentation, among others. And we are privileged really to be helping some clients already to develop very tailored toolkits that they use to assess their AI systems to comply with all that.

Julian Dibbell

Okay, so that's complex work that needs to be done on the governance front. How about technology transactions, deals? How are those going to be impacted by this new legislation?

Ana Bruder

I think that the UAI Act will actually help businesses that are purchasing AI, the customers. Because the UAI Act will clearly stipulate the obligations that fall upon providers of those systems and which I just mentioned is actually the majority of the obligations, right? So that means that the terms and conditions of the providers of AI systems will need to be amended to reflect that. And technical documentation will also need to be provided to the customer by the providers of the system, ideally also attached to the contract so that the employer of the system can follow any applicable instructions. This is written in the UAI Act quite explicitly. So customers, the deployers of the AI system will have an interest in making sure that the terms of the provider are amended because that will alleviate some of the obligations very likely currently falling upon them. Or let's say, missing reps and warranties given by the provider. What else? Well, actually the provider will actually also have an interest in making sure that the contracts reflect the obligations that fall upon the deployer of the system. For example, if the customer is just purchasing an AI model, but the customer is the one which will be the deployer, right, inputting data into the system, getting the output data. And so all of the decisions around data governance will actually be in the control of the deployer. And so those obligations relating to data governance, that the data be accurate,

adequate, et cetera, those will fall upon the deployer and that should also be reflected in the agreements basically. So we will definitely see the purchase of AI systems undergo some contractual changes with the adoption of the UAI Act.

Julian Dibbell

All right, so it sounds like there's a lot of work to do here, both on the internal governance front and also on the contracting front, both reviewing contracting practices, also existing contracts. What is the timeline for all this? I mean, in terms of that staggered set of dates you talked about, I mean, is this something that all has to be done within the next six months or is it going to be tied to those rolling adoption dates?

Ana Bruder

So for Gen AI, the Gen AI provisions, we'll start applying within a year, right? So that's actually not a long time, if you think about it.

Julian Dibbell

Generative AI you're talking about, these new large-language models and so forth.

Ana Bruder

Right, so the UAI Act uses a terminology. Yeah, exactly. So the UAI Act calls them general-purpose AI, but that's what we're talking about really. So the providers of those models, they really need to start thinking about amending their terms right away, because one year after adoption, that will apply already, right? For high risk, if it's a high-risk AI system, like the one, the example I gave for the employment tool, and there are many tools like that already on the market, that will be with, you know, two years, or in some cases, depending on the specific-use case, could be three years after adoption. So there is a little bit more time.

Um, yeah, I don't think you would need to start thinking about this right now, but you know, as Arsen will talk to us about in a minute, there are other considerations that you need to make internally, right? Regarding AI governance more broadly, you know, even thinking apart from the tech transaction side of it.

Julian Dibbell

All right, well, a lot going on in the EU. I want to turn now to Oliver for the update from the UK. What's going on there?

Oliver Yaros

Thanks Julian. Well, it's a pleasure to be talking to you. There is quite a lot going on in the UK. I think the UK, it's fair to say, is taking a bit of a different approach to the EU approach. The UK government about a year ago in March announced in its White Paper that it wanted to take what it called a pro-innovation, principles-led approach. And that is that there wouldn't be a new regulator created for AI or new regulations specific for AI, but that existing sectoral regulators would be empowered to make additional guidance, which would explain to providers and employers of AI technology, what they needed to do when deploying an AI system for use in the UK.

And that approach was confirmed a couple of months ago in February in a response to the UK government to a consultation process that it had led since the White Paper last year. So the idea is that

there's going to be a central function within the government that will explain what the approach to be led by the UK government is going to be.

There'll be someone who's representing AI interest in all government departments and an expansion of the AI team within the Department for Science, Innovation and Technology or DSIT as it's called. And then the existing regulators will together form something called the Digital Regulation Corporation Forum to coordinate the giving of advice on the use of AI and the main regulators that I think our clients will be most interested in are the ICO which is the data protection regulator which regulates the broadcasting and communications, the FCA which is the financial services regulator and the CMA which is the competition and markets authority and they will be issuing guidance and have already issued some guidance around five principles that the UK government wants organizations to think about when deploying and creating AI technologies. Those principles are that AI technology should be safe and secure and sufficiently robust. That's the first principle. That there should be sufficient transparency and explainability when using AI, that the use of it should be fair, that there should be accountability and governance in terms of how it is used, and there should be a right to redress and to contest decisions that are made using AI technology. And these broadly correspond to the OECD values-based AI principles that we've seen other countries focus their efforts around internationally.

So what sort of guidance is already out there and is going to be issued? Well, what I think is quite interesting is that DSIT has published guidance to support how the different regulators should think about issuing guidance on AI regulation going forward. It's got some non-binding suggestions for regulators to follow and they discuss how each principle can be adopted in turn. The ICO has actually produced quite a bit of guidance already. For example, if you're using AI systems that are going to involve the use of personal data, there is guidance by the ICO that talks about how you should do your data protection impact assessment to assess the effect that using AI technology will have on people's privacy rights. They've also produced two other pieces of guidance: how to explain decisions made using AI and how to use AI with biometric technology. The Competition Markets Authority, which is mainly focused on encouraging competition within the UK, has issued some interesting guidance and a report on foundational models and how to make sure that the supply and use of these models remains competitive in the UK market.

So we're going to see some further guidance issued over the next few months towards the end of the year. We're going to see a cross-economy AI risks register being prepared by DSIT over the course of this year. We're going to see some further guidance on how you limit the different purposes for which generative AI may be used in terms of the life cycle of technology, later this year we may also see some further guidance on the use of AI and recruitment and HR. And of course, most interestingly, we're going to have an election later in the year in the UK, so we can potentially see a change in approach there as well. One final thing I wanted to mention is that the UK and the US recently signed a memorandum of understanding to work together to develop tests for the most advanced AI models and that really builds on the AI Safety Summit which was hosted here in the UK towards the end of last year and is an agreement between the UK and US AI Safety Institutes about how they can jointly focus on promoting safety with the use of AI systems going forward. So that's a bit of a view as to what's happening in the UK.

Julian Dibbell

Well, in contrast to the EU situation where we have a single piece of legislation kind of dominating the

conversation, this is a bit more of a diffuse and emerging landscape, it sounds like. I'm wondering, you know, in practical terms, what does this mean for organizations with a presence in the UK? How can they start taking steps to address the existing requirements and perhaps the ones coming down the pike?

Oliver Yaros

I certainly think you're right, Julian. This is a bit of a different approach. I think first of all, clearly the UK is an important market and it's important for clients to be thinking about how AI technologies are going to be used in the UK. Clearly most organizations will want to be taking a similar, if not the same approach, across their organization internationally and I think Arsen will talk a little while about how organizations should be thinking about doing that. But with respect to the UK specifically, I think the important thing is to think quite carefully about what AI technology might be deployed in the UK and how it's going to be used and then what guidance, how compliance with specific guidance related to that technology, can be complied with when deploying the AI technology. So for example, the ICO guidance on how to do an impact assessment where AI uses personal data is something that clients should really be thinking about when using any AI technology in the UK, which is going to involve the use of personal data. Obviously, there's a whole weight of things they have to do under existing privacy laws like the GDPR and that applies across the EU and the UK, but there's additional guidance they need to think about complying with when assessing the risks of doing that when using AI technologies in the UK. In particular, when you're using AI technologies to make decisions in the UK, and it might be in an HR context, or it might be in another context like credit decision-making, I think it's important to look at specific guidance that the ICO has produced in that context. So I think when organizations are thinking about how to use AI in the UK specifically, they should be thinking, they should be looking to define quite carefully what the use-case is, think about what types of data points are going to be used when using that AI technology, and then thinking about what guidance might be relevant that they need to take account of when doing that before deploying the technology in the UK. So those are my thoughts there.

Julian Dibbell

All right, thank you, Oliver. Arsen, I want to turn to you over here in the US and get your perspective on the US and global landscape for laws and regulations governing artificial intelligence.

Arsen Kourinian

Great, thanks Julian. So let me just start with the worldwide approach, which sort of provides the baseline and then go into the US. There's different approaches different countries are taking with respect to managing AI risks and implementing, requiring companies to implement governance. One is the most flexible approach, which is the principles-based approach. This largely stems from the OECD's AI principles, which countries on six continents have adopted and including the G7. So as a baseline majority of the countries worldwide that are adopting AI legislation or approaches or government guidance, they're really formed based on the AI principles from the OACD. Looking at those principles, then we shift to the different approaches of how to implement those principles.

A country like Singapore, for example, has decided to take a voluntary approach where they essentially gave the tools companies need to implement some of these principles, such as a toolkit that they have, along with a guidance that they've issued about how to implement proper AI governance within your organization. And then other countries have taken varying approaches as to how to implement it strictly from a sector- or context-specific approach, which Oliver talked about, which is happening in the UK.

Same is happening in the US on a federal level, where essentially the position of these governments is that we do already have enough laws on the books to address AI. Here are some guiding principles and we just need our regulators to enforce and make sure that companies are using AI in a safe, secure and trustworthy manner.

In addition to that, the US has had some of the major tech companies voluntarily commit to certain principles that happened last year, and then they've also issued guidance in April of last year about the various agencies in the US that are going to enforce existing laws.

And then there's the more role-specific and risk-based approach that some countries are taking. We heard from Ana about the EU AI Act. That's probably the gold standard at the moment where it demarcates obligations based on the level of risk and the role that you occupy within the AI ecosystem. Canada and Australia appear to be headed in that direction as well where they're going to have some level of a risk analysis and role-specific context as to how companies need to implement AI risk-mitigation measures. And then interestingly, the US in general, as I mentioned on a federal level, although it does appear to be taking a sector- and context-specific approach similar to the UK; on a state level, however, we're seeing a number of bills that are essentially many versions of the EU AI Act, where they have various obligations that are tied to risk, high-risk type of activities, and also the obligations that you're subject to are contingent upon whether you're a deployer of AI or a developer of AI. And so until we get federal legislation, it's possible that we may see, similar to our US privacy laws, state-by-state various AI laws passed that are comprehensive in scope and similar to the EU AI Act.

Julian Dibbell

Alright. Well, a lot of varying approaches here and changes coming down the road. I'm going to ask the same question I asked Oliver. In light of all of this variety and change, what are the practical steps companies can take to implement compliance with all of these directions and guidance and regulations?

Arsen Kourinian

I think for starters, you want to have an appropriate infrastructure in your company to be able to address all of these requirements. And I think if you try to address the various, especially for multinational companies, if you try to address the various different approaches some countries are taking, to start off with a checklist of all these nuanced issues that a particular law provides for as the initial starting point as opposed to part of the steps. It could be challenging. It's a bit of playing Whack-A-Mole where different requirements are springing up, and you're sort of addressing them ad hoc without an adequate strategy. And that could prove challenging in case there are some inconsistencies or different jurisdiction approaches. So what I like to recommend to companies is start by actually implementing your holistic AI governance program.

And so the good news is that it's possible to do this on a global scale because the trends we're seeing, at least they're based on common components for compliance and then just the implementation aspect of it may require a couple of checklists of issues. So let's jump into that. What does that mean? What does AI governance mean? Well, for starters, you need an AI governance team. You need a team of multi-stakeholder members that are skilled in different skill sets and are able to address different components of AI governance such as IP, data privacy, confidentiality, HR, marketing, procurement, you name it. And so once you assemble an AI governance team which has varying and diverse skill sets, those are the individuals that are going to be your sort of ethics board or AI oversight board that's going to give the

top-to-bottom direction for how to implement compliance for companies that have a decentralized approach for management. You should still think about a bottom-up reporting, even if you delegate compliance on a local country-by-country level, so that there's an oversight board that considers what's going on in different geographic regions and can give further guidance. Next, another important component is data governance.

So if you're a developer of AI, where you're making the AI systems, you need to make sure you're training the AI model or fine-tuning a foundation model using high-quality data that is properly annotated, is representative of the environment you're in, has proper rights to the data that is being used to train or fine-tune the model. If you have the appropriate privacy rights, licensing rights to various data or IP rights. So you need to make sure you have the proper rights and high-quality data to train and develop your model. Next, you need a risk management plan. And so...

Ana talked a little bit about what the EU AI Act requires, but basically there's a risk-ranking of prohibited, high, limited-to-minimal risk. That's generally the approach a lot of other countries are taking who have adopted a risk-based approach. You need to document the level of risk, the type of mitigation measures you're taking, the benefits of using AI, the probability and likelihood of harm, and document all of this in an AI-impact assessment. And so for any data privacy professionals out there we're used to doing this, but this is now broader in scope because we're not just dealing with personal data, but broadly any data and any functionality of an AI to document this process. The next step is the Whack-A-Mole aspect of it, which is the legal compliance. As you can see, different countries may have minor and broader various privacy laws that trigger AI, IP laws that trigger AI, AI-specific laws, like Utah passed a minor AI law related to Gen AI, being transparent about it. So you need to be able to identify how do I address this as part of my governance program. And so that's where your local counterparts come into play where you would do your typical compliance oversight of there is a law, here are the requirements, how does this fit into our governance program, have we addressed all of these or not, identify the delta, and then bridge that gap. Next, what AI governance is, you need to implement the full scope of mitigation measures. There are a number of them. You need to be transparent about your practices. You need to explain how AI works. If you're a developer of AI, you need to give instructions of use to the deployer on how to use the AI, the limitations of the tool, various technical information. You need to also do some testing to make sure that your AI is fair and unbiased using appropriate methodologies. And then you need to test the accuracy of your AI model. And there's different ways of going about it.

You also need to continuously monitor AI. It's not a one-and-done project to implement this governance process. Rather, it needs to be done throughout its entire life cycle, and you need to go back and make corrections because AI does drift and it does hallucinate with Gen AI, and you need to be able to identify what the issues are.

And so finally, the last step is accountability where you need to document through appropriate policies and procedures like an internal use policy or an AI development policy, how to have this full AI governance.

Julian Dibbell

Alright, well, thank you, Arsen. Thank you, Oliver and Ana. A lot to chew on, as is the case with AI always. Thanks again.

Listeners, if you have any questions about today's episode or if you have an idea for an episode or you'd like to hear about anything related to technology and IP transactions and the law, please email us at techtransactions@mayerbrown.com. Thanks for listening.

Announcer

We hope you enjoyed this program. You can subscribe on all major podcasting platforms. To learn about other Mayer Brown audio programming, visit mayerbrown.com/podcasts. Thanks for listening.