

Podcast Transcript - Tackling the Complexities of AI Contracting: Practical Guidance and Insights

In our latest podcast episode from Bird & Bird, Will Bryson (Partner UK, Bird & Bird), Emma Ren (Associate China, Lawjay Partners), Clemens Molle (Associate Netherlands, Bird & Bird) and Omar Sharief (Associate UAE, Bird & Bird) share their insights on the most common challenges, and hidden pitfalls when contracting for AI systems.

July 2025

Emma Ren:

Hello everyone. Welcome to this Bird & Bird podcast. I'm Emma from our Shanghai office and I'm delighted to be joined today by Will in London, Clemens in The Hague and Omar in Dubai. In this episode, we will be exploring some of the most common challenges that companies face when procuring or selling AI systems and sharing insights from our practical experience in this fast evolving area.

And to give you all a little spoiler, we will be hosting a webinar this autumn for a deeper dive into contracting for AI issues. And we will also be launching the second edition of our Contracting for AI Toolkit at the same time. The link to the toolkit is in the show notes and is free to access. So be sure to check our websites for updates, join the upcoming webinar and take advantage of our toolkit to help you navigate these challenges in your own work.

So to start us off, Will, what do you think are the most common oversights or misconceptions that Buyers and Sellers may have when it comes to procuring or selling AI systems?

Will Bryson:

Thanks Emma. Look, I think one of the traps we see quite often is businesses thinking, "Well, AI is just software. And we've been negotiating contracts for software for years. So we're good to

go. We don't need to do anything different here." But fundamentally, the intricacies of the technology and how it functions, how it's put together, together with this emerging and developing regulatory landscape that we have means there are areas that do need to be considered differently when you're thinking about how to approach a contract for AI systems.

Clemens Molle:

So maybe just to jump in for IP. I think one of the hottest topics is whether AI output can be copyright protected. So up until now, judgments throughout the world have been quite mixed. But there have been more judgments where copyright protection was denied rather than confirmed. But regardless of whether output is copyright protected, parties can still make contractual arrangements about the use of output. So parties still need to consider this in any contract. And also, should the output be copyright protected, we see that buyers are sometimes wondering if they will become the owner of those rights. And I think it would be a misconception to think that sellers never agree to give a buyer these rights because they often do. And with generative AI, sellers on the other hand might assume they only need to provide limited usage rights to the buyer, or sellers might assume that they can use the inputs of the buyer to further train their AI system, also for other customers. So the contract needs to be clear on these points to avoid disputes.

And finally, another topic we see in practice is that AI systems sometimes rely on open source software or licensed third-party components. And contracts may neglect how these components are used and that potentially exposes the buyer to restrictions that the open source licenses impose.

Emma Ren:

Yeah, looks like IP ownership is really a difficult question. Omar, what's your view?

Omar Sharief:

Yeah, look, when it comes to AI procurement, I'll be speaking from a UAE perspective. And for the UAE, one misconception that a lot of people tend to have is that it's treated as one monolith of a jurisdiction. But actually, the UAE itself is a federation of seven Emirates, where each Emirate has their own Emirate specific laws. And then there's also federal specific laws that affect the UAE mainland. And just to complicate things a little further, there are also things called free zones. There's more than 40, but I'm going to go over just two main ones, which are the Dubai International Financial Centre, aptly called the DIFC. And the Abu Dhabi Global Market, which is known as ADGM. What's unique about these two free zones in particular is they actually have their own court systems. They actually even follow common law, largely inspired by English law. So a buyer might assume that a contract for Dubai will work in the same way in Abu Dhabi, but that's actually not the case. In fact, a contract in Dubai, well, is it the mainland? Is it a free zone? And then even in Abu Dhabi, is it the mainland? Is it a free zone? And each one has different laws. So for example, in the UAE, you've got three separate data protection regimes. You've got the federal law, which applies across all the Emirates in the mainland. Then you've also got DIFC's own regulations, and then ADGM's own regulations. And each one has their own requirements for AI systems, for automated decision making, and other matters like that. So if you're deploying across multiple Emirates or free zones, you need multiple compliance strategies, not just one.

And then a second point is between laws and guidance. So in certain jurisdictions, for example, in the EU, you've got the AI Act, which is quite hard law. But the UAE actually operates a bit differently in that there isn't one comprehensive law. And in fact, it's actually more guidelines and the AI strategy, ethics guidelines, AI charter, sector-specific guidance, specifically when it comes to AI. So a seller might make the mistake of having a Middle East-compliant solution. But actually, you couldn't even have a UAE-specific

solution, frankly, because there's just so many factors involved with understanding what is required. And you might have UAE vendors demonstrate specific ethical principles and national objectives. So it's not something you can just tick off as a compliance.

That brings me to the next point, which is cultural considerations. That can be a major blind spot because various AI systems are actually trained on the English language, on Western data sets. But in the Middle East, that doesn't actually reflect the local context. So for example, Arabic language capabilities, or even to some extent, Islamic values in the AI design. And sellers will actually frequently underestimate how important that localisation is. And it's not just about translation. It's about understanding local cultural nuances. So someone from Iraq and someone from Tunisia would have vastly different cultural nuances, which is very important to understand, particularly for customer-facing AI systems.

And then finally, I'll just touch on the evolution of AI in terms of in the UAE. UAE regulations are and have been deliberately made flexible. And that's why their guidelines are not laws. Because there's an appreciation that they're going to change over time. And there's anticipation that it will happen very soon. And so from the UAE perspective, laws can remain static, but guidelines can be adapted. But it's important from a buyer and seller point of view that to include regulatory change of management clauses, regular compliance reviews. AI legislation is constantly going to change so flexibility is essential.

Emma Ren:

I can really see that. People really like flexibility, but that brings in more uncertainties. And we need to manage that in the contract, right? And I also echo from the Chinese perspective that the difficult and different jurisdictional principles and regulations need to be taken care of when drafting a contract from international scope of view. Thank you very much for your answers.

So for the next question, what are the clauses that you find most commonly attract disputes? How have you handled them in your real work, in your previous practices? Clemens, would you like to start?

Clemens Molle:

Sure. So like Will already mentioned, AI is really quite different from other pieces of software. So also when contracting, you need to take that into account. And I think one of the instances where

that really comes to light is that AI systems can sometimes evolve and improve over time. Sometimes these improvements happen because of the interaction of the buyer. They're with the system, either through inputs or follow-up questions or corrections of the buyer. And in a sense, by doing so, the buyer further trains the AI system with its data. And if a contract does not clarify who owns such improvements or how they can be used and by whom, disputes can arise. So sellers might aim to retain ownership of the core model plus improvements, while buyers on the other end may also want the rights to the improvements. And I think that's especially true if the buyer contributed its proprietary data. But the issue is, these improvements to the AI system are often quite difficult to separate from the model or cannot be used without the model.

Will Bryson:

Clemens, just to chip in there. I completely agree. That's something that we see all the time. It's this tension that you can't really separate the trained model from the base model. And what does it mean in practice to say that you own the new layer? If you can't disaggregate that or apply it elsewhere. And I think that really calls into question, what is the value of owning that new IP if you can't really separate it? And I think the answer to that question is probably a defensive one. I think customers are trying to get in a position where they are protected from the supplier itself using a model that has been customised for them. And it might perform in a way that they don't want it to for another person, particularly if it might be a customer of theirs. But you don't necessarily need ownership rights to achieve that. And what we're seeing quite a lot is the same effective position being reached through exclusive licenses and then with an obligation to erase or destroy the system at the end of the contract. So if it was a SaaS system, for example, an undertaking on the supplier that though they may own the trained model and that is exclusively licensed to the customer, the supplier will undertake to destroy that model on termination. So the supplier just left with that base model that they're obviously free to reuse for any other customers.

And also just to chip in on one point, Clemens, you mentioned earlier on data rights and the ability for suppliers to reuse the data the customers provide to them. I totally agree that's an issue that people often overlook and they need to think about how that's dealt with in the contract. But I also think it's quite often an area of contractual dispute. Perhaps less so these days, but in the early days of generative AI, there was

obviously a lot of appetite from suppliers to kind of reuse that data themselves. And obviously that creates tensions with customers who don't want any possibility of that potentially sensitive or proprietary data popping up elsewhere, even if that isn't likely in practice. So that's often an area of some negotiation. I think we are seeing a move from a world where the suppliers tend to be asking for that to one where typically they're less concerned. They've got enough data from other sources, abilities to synthesise their own relevant training data, that actually trying to Hoover as much data up from customers as possible is less important, but certainly can be a contentious issue.

Emma Ren:

Clemens, do you think there are any other points of contentions that you have come across?

Clemens Molle:

I think so. What I've seen is that AI systems can, of course, produce output that may infringe third-party IP rights, or AI systems can rely on training data that was used in an infringing manner. So that begs the question of who will be responsible for this. Buyers, for example, may assume that sellers will provide quite broad indemnities for IP infringement claims, while sellers might try to disclaim or limit such liability. And I often see discussions about infringement indemnities in AI contracts. So a contract should define each party's obligations and remedies if an infringement claim arises. I think in practice, this is often clear for the training data, but often less clear for potentially infringing output of the system.

Will Bryson:

Yeah, I would certainly echo that, Clemens, because I think in what I see, that is often a bone of contention. Certainly, I think in the early days of generative AI particularly, there was a reticence from suppliers to give any form of indemnity protection in relation to outputs. There was so much unknown about how it worked, what the risk was, that it didn't really release something that was offered. I think these days, kind of having a negotiation around an indemnity for infringing out, but it did be much more common, and suppliers are willing to give those. But I think we often do expect to see our carve-outs to those indemnities, which kind of go beyond what you would typically expect in a traditional software contract. So it's not just, combination with a third-party piece of software or modification or use in breaching of the agreement. It is things around guardrails in place, deliberately trying to elicit infringing material. If the

training data or the input was the thing that caused the output to be infringing, these sort of very specific exclusions and carve-outs from traditional IPR indemnities, I think, are quite commonly asked for by suppliers and then tend to be points of negotiation.

I think, in my experience, an area of contention is around explainability and accuracy of outputs. So customers, understandably, seek assurances from their suppliers that what the system produces is accurate, that it's correct, that they can rely on it and they can use it. But they also want to know that an explanation as to how that conclusion was reached, is the output explainable? Can I know how from a given input I got to a given output? The issue here is these are probabilistic systems. They're not deterministic. You can't draw a straight line between input and output and say exactly why these happened all the time. And neither can you, because of that, necessarily say that an output is going to be accurate. It is a, largely probabilistic outcome. So, it's probably right. It's got statistics behind it, but you can't say it's right in the same way. And that becomes a tension within contracts. Customer asking for X and supplier saying, "I can't, I can give you Y." And what that Y looks like will depend obviously on the systems themselves and the negotiations as they go. But we do see a lot of compromise positions around using good industry practice. The training system will endeavour to produce accurate outputs. We can see that, but also people do understand that things like hallucinations are just part of the technology at the moment. And just have to be built in and factored into how you're going to use the outputs. If you can't rely on the accuracy, well, that can be okay. You just need to think about how does that factor into how I'm using this output? Do I have a human in the loop? Does it matter if it's right? There's lots of different solutions to this particular problem.

Emma Ren:

I can see the difficulty in that. Even if you input the same word in the ASMA model, it may give you a different answer every time. Thank you, Clemens and Will for these insights. I think the ownership of improvements, the infringement indemnity clauses that you've mentioned, and the challenges around proving accuracy or sustainability are all the common points of contention that companies can expect to face when they are negotiating for AI contracts. It's clear that without specific contractual terms, both buyers and sellers can find themselves in tricky situations.

And with all these potential issues in mind, I'd like to move to our next topic. Do you think it's

worthwhile for businesses to prepare a template AI contract or even just a checklist before they dive into any AI project? What have you seen in practice? Is there any real benefits to doing so? Or maybe if there is any drawbacks? Clemens, what's your view?

Clemens Molle:

Yes. Well, I do think it can be helpful to draw up a checklist in advance, or at least think about a couple of points before you enter into negotiations for a contract. So a buyer, for example, can think about if it wants to allow the AI system to be trained on that data. And if they do want that, who can then reap the benefits of that improved system? And for sellers, this point is also relevant, as is the point in how far a seller wants to go in terms of providing the indemnities that we talked about for third-party infringement claims. And of course, this checklist and this template AI contract mainly relates to the bigger, individually negotiated contracts. But for smaller companies who do not individually negotiate these contracts, it can also be helpful to think about these points. General terms and conditions can vary quite a bit between different AI providers. So in that sense, you can also pick and choose, so to say, without individually negotiating a contract.

Emma Ren:

Thanks, Clemens. Omar, Will, any points to add?

Will Bryson:

I think I would just say, I think it's absolutely a worthwhile exercise. Both as Clemens says to think about the issues and produce a checklist, but actually to go that next step to actually have a set of terms and conditions that you are ideally able to use. Now, whether or not you're going to be able to use your own terms and conditions as a customer in particular in any given circumstance will obviously depend on the particular deal, the bargaining position, what you're buying, if it's off the shelf, if it's customised. But even if you're likely just to be largely buying off the shelf, non-customised stuff, I think it's still worthwhile. It really allows you, as Clemens mentioned, to think about the issues, to go through each of the topics, each of the risk points that may arise in the system and work out what your best case scenario is and what your negotiating parameters are. And to do that in an informed way. Then you, as a lawyer, can go into a negotiation and say, well, I know what the issue is here. I know why this is a problem. I'm not going to waste time

trying to push for positions that are just simply unobtainable, whether that's asking for 100% accurate all the time sort of warranty where a system just can't do that. And you don't miss things, like we talked about ownership of outputs and if you can't own outputs, how does that work? The fact you've thought through everything, I think is just invaluable. And it's, I think, a worthwhile exercise, regardless of whether or not you will always be able to rely on your paper. I think that always comes with a pinch of salt. You need to obviously know the times when there's no point even deploying your paper. If you are literally just buying a few off the shelf licenses for an enterprise product there's probably going to be limited scope to negotiate. But again, at least you're informed, at least you can have that conversation with the business to say, well, look, these are the risk points. This is how the contract, their contract addresses it. This is how we're going to mitigate those risks. So obviously, I would say this as a commercial contracts lawyer, but I genuinely think it is an important thing for any business who's going to be procuring AI systems to start thinking about doing.

Emma Ren:

It looks like taking some time to prepare can really make a difference. Omar, what's your point?

Omar Sharief:

I think from a UAE perspective, given the split legal landscape, a checklist is probably more valuable than a template. And I say that because there are so many different laws and regulations to factor in, that you'd effectively just be drafting various templates that it probably would lose its value in that respect. But a checklist certainly would be valuable to ensure that you've factored in various obligations on data usage or automated decision making, things like that. And I also think that because of the guidelines, they're not laws, the guidelines, they're always changing, they might expand, they might shift, things like Arabic language content, local cultural values, I think those are important to include in a checklist, just to make sure that the checklist covers those important aspects. You could maybe arrange a template between Dubai and Abu Dhabi, particularly DIFC and ADGM, which are both common law jurisdictions. And you know, the laws can be regarded as a little bit similar, but there are still differences that I think a template would still need to be changed between each.

Well, as Will said, I think it's important to know from your perspective, from a negotiation point of

view, you know what your position is around data protection, local compliance, and just having a change of law or change of regulation, clause can be really important because there might be AI specific requirements down the road.

Emma Ren:

Yeah, I also really think it's quite beneficial to have the company's own position beforehand, and they build on to their different contract clauses based on that.

Thank you very much, Clemens, Will and Omar for sharing your expertise and practical insights today.

For our listeners, I'd like to mention again that we will be exploring these topics further at our autumn seminar online. So we hope you will join us virtually October. See you around at that time.

twobirds.com

Abu Dhabi • Amsterdam • Beijing • Bratislava • Brussels • Budapest • Casablanca • Copenhagen • Dubai
• Dublin • Dusseldorf • Frankfurt • The Hague • Hamburg • Helsinki • Hong Kong • London • Lyon
• Madrid • Milan • Munich • Paris • Prague • Rome • San Francisco • Shanghai • Shenzhen • Singapore
• Stockholm • Sydney • Tokyo • Warsaw

The information given in this document concerning technical legal or professional subject matter is for guidance only and does not constitute legal or professional advice. Always consult a suitably qualified lawyer on any specific legal problem or matter. Bird & Bird assumes no responsibility for such information contained in this document and disclaims all liability in respect of such information.

This document is confidential. Bird & Bird is, unless otherwise stated, the owner of copyright of this document and its contents. No part of this document may be published, distributed, extracted, re-utilised, or reproduced in any material form.

Bird & Bird is an international legal practice comprising Bird & Bird LLP and its affiliated and associated businesses.

Bird & Bird LLP is a limited liability partnership, registered in England and Wales with registered number OC340318 and is authorised and regulated by the Solicitors Regulation Authority (SRA) with SRA ID497264. Its registered office and principal place of business is at 12 New Fetter Lane, London EC4A 1JP. A list of members of Bird & Bird LLP and of any non-members who are designated as partners, and of their respective professional qualifications, is open to inspection at that address.