TECH REIMAGINED
AN ENDAVA PODCAST

AI & The Law

[00:00:11]
BRADLEY HOWARD, ENDAVA HOST (BH): Hello, I'm Bradley Howard, and I'm happy to welcome you back to the latest episode of Tech Reimagined. It's my pleasure to introduce you to our guest today, John Byers, who's a partner and commercial solicitor at Osborne Clark. Hello, John, how are you today?

[00:00:26]
JOHN BUYERS (JB): I'm fine. Thanks. Bradley, thanks very much for having me.

[00:00:33]
BH: So the topic of today's episode is A.I. That's artificial intelligence from a legal perspective. So John, can you tell us a bit more about yourself and your background and artificial intelligence as well?

[00:00:40]
JB: As you've mentioned I’m a partner at Osbourne Clark, which is an international law firm and I've been in the I.T. industry as a lawyer for more years than I care to remember. I started off as an outsourcing lawyer working on traditional outsourcing deals, then moved progressively more and more into digital technologies. Actually, it was the evolution of the outsourcing industry itself, which kind of pulled me into artificial intelligence, and it's an absolutely fascinating topic. The upshot is that I now lead the international team at Osbourne Clark, that advises on artificial intelligence and machine learning. I’m also in an organisation called the International Technology Law Association as a board member, and I'm also responsible for their A.I. committee. I've written a book on the legal implications of artificial intelligence, which is now in its second edition. So I'll have to get a second edition to you as well, and there's a third edition in the pipeline. With my AI Tech Law hat on, we've also published a comprehensive guide to machine learning and artificial intelligence ethics. So, yeah, that's me really.

[00:01:54]
BH: Well thank you very much, John, and it's amazing to have such an expert on the show today. Typically, when people think of AI, they think of something out of a science fiction movie. But of course, it's completely in our day to day lives through navigation apps or ride sharing, video apps that recommend the next video or email spam filters. The list just goes on and on, so it makes sense, of course, that there should be a legal aspect to it which you were discussing in your introduction as well. So can you just walk us through some of the legal issues that software developers might need to think about when they're developing A.I. platforms?

[00:02:31]
JB: Yeah, sure, I'll do that, and it probably makes sense for me to just give a brief overview of the different types of A.I. technologies, because I think there's a lot of confusion in the market about what artificial intelligence actually is and a lot of frankly ill-informed journalistic hyperbole. I don't want to do my journalist colleagues down, but it's quite a catchy strap line to talk about artificial intelligence, machine learning, and inevitably, you end up with, as you've been intimating, one of two binary outcomes. Either it's the end of the world with the dystopian robots or the beginning of the future where AI is going to save us from mutually assured destruction. I think the reality, as you've indicated, is somewhat more prosaic, it’s quite an exciting technology, but it's very specifically directed. We're nowhere near what some commentators have called the singularity, which is when machines become self-aware.

What we have now is a situation where we've got some very, very clever, specifically directed applications. Let me kind of unpeel the onion. If the listeners can kind of have a metaphorical onion in their head, when I'm when I'm talking about it, the way I kind of visualize artificial intelligence is you have the outer layer, which is A.I., which is the ensemble term, which covers a variety of different technologies, some of which are a black box and we'll get to that in a minute, and some of which are not. So you might have something called a decision tree, which is an assistive AI technology, which essentially conceptualizes that as an automated yes or no flowchart. So that would help an AI solution, and underneath the AI layer, you've got machine learning, and that's the really exciting area that all of these legal issues stem from.

This is not a new technology. This was something that was conceptualized back in 1952 by a chap called Marvin Minsky, who wrote about it with another guy called Seymour Papert, and they just didn't have the technology infrastructure or the scale of computing power to actually make it useful. That's where the neural network sets, which is a form of artificial intelligence that seeks to mimic the way the human brain works. Finally, within the machine learning layer, right at the very core of the onion is where the really exciting stuff is, which is deep neural networks, and these are machine learning neural networks that are scaled up using the power of modern technology. This is really where the kind of secret sauce of machine learning actually is. So that's a kind of explanation of the different types of technology that together come to form artificial intelligence.

In terms of why it's different and how it's different from traditional computing systems, it's probably worthwhile explaining that, you know, we're all aware and illustrating the issue that causes a lot of the legal the downstream legal problems. But we're all aware that a traditional computer system is essentially a computer that operates a prescriptive program, and it spits out a defined solution. That is something that we've been very used to for the last 60 or 70 years. That's traditional computing. You have a human operator that types a program and produces a result. Very good at lots and lots of routine transactions, not very good at independent decision making or coming to its own conclusions. It's a very unforgiving environment.

If you get a computer program wrong, it will spit out a syntax error. Contrast that with a machine learning situation, and a machine learning solution. What actually happens there is that you assemble, you have a problem and you assemble a dataset, a collection of data that is used to train a machine learning system to come up with a solution to the problem, so you're not - there is an algorithm which defines the way in which the machine learning system reviews that data, but you're not writing a prescriptive program for it. It's coming to its own conclusions. Therein lies the problem, because what happens is a machine learning system will come up with an answer, which will be on a variable scale. So just as you, Bradley, if you come up, if you make a decision in relation to something, you're going to come to a conclusion on a range of probabilities. A machine learning solution will do exactly the same thing, and the important point to note is that these decisions at the margins of a probability curve.

What the technologists call edge case decisions. They're not faults. They are within the bounds of probability. They're just less usual. This causes no end of problems for lawyers because we're used to very clear concepts of cause and effect. We like to know who is responsible for a problem, and the law is geared up in that way to provide a remedy to someone who's been harmed. If you can't identify a particular fault, then you can't serve justice or provide a remedy to a person that's been injured. That's fundamentally at the heart of machine learning. I ought to kind of elaborate in that. There's not only the kind of case decision problem, which is the kind of variable solution issue which taxes lawyers, but also at the heart of neural networks, we have a fundamental technological conundrum, which is that black box solution.

So I alluded to that earlier in my answer. What that means is that even though data scientists can create the algorithms within which these systems operate, they simply don't understand how or why the machine learning system, particularly in the context of deep neural networks that have lots and lots of processing layers, come to the conclusion they did in a particular circumstance. It's not explainable, and that causes all sorts of issues, and it compounds the causation issue that I've just discussed in terms of determining who is responsible for a fault. So that's a fundamental break point in relation to serving a remedy up for somebody who's been harmed.

[00:08:54]
BH: So when someone comes up with an original idea, at what point would you like to be engaged?

[00:09:00]
JB: The answer is that we need to be involved as soon as possible, recognising that lawyers cost money. But I think we do need to educate the client to a certain degree on the particular part of their solution that will cause legal issues. So we talked about the causation issue with the use of neural networks, and that's really not something that can be remedied particularly well. It’s there, it’s a feature of the technology. So I guess if you're a client that is about to use a machine learning solution and we're going to get to regulation later. But if you're about to use a machine, you need to ask yourself whether it is worth the inevitable regulatory risk to employ something like artificial intelligence, or whether it can be dealt with by a more prosaic technology, a more conventional technology.

That's certainly a decision gate that lots of businesses are going through at the moment. What you simply can't do is implement artificial intelligence because it's the next bright, shiny thing and it will do it. There are regulatory hurdles that we'll talk about in a minute, which will cause you to pause and think and pitfalls that you fall into. As a developer, one of the other issues that need to bear in mind is, well, there are loads! But let's start with bias and how machine learning solutions can inadvertently create bias in outputs. This is to do with how you select your datasets to train the machine learning systems. One of the most interesting examples that I've come across, which really brought the whole issue to light to me was listening to the head of Google Translate. He was talking about how Google developed their Translate algorithm, which is one of the largest, has one of the largest datasets on the planet actually. As you can imagine, it's taking in all sorts of data sources around the world to create a multi-language translation platform. What Google found when they deployed Google Translate was that the algorithm inadvertently conveyed the prejudices of the texts that it had been reviewing, and Google had given it lots and lots of texts that were out of copyright. Very old texts like the Bible.

These outmoded cultural positions were reflecting its outputs. There's a very interesting example of a Google Translate, taking English and translating into Finnish and then back into English and likewise with Turkish - which are less… I’m not an expert in Finnish or Turkish, but they're less gender specific. So if you said ‘he is a nurse’, Google Translate would come back with ‘she is a nurse’. If you said or typed in, ‘she is a doctor’ it would come back with, ‘he is a doctor’. If you said ‘he is lazy,’ it would come back with, ‘she is lazy’, so clearly completely unacceptable from the point of view of Google to have that kind of inbuilt stereotyping and outmoded cultural attitudes. What Google had to do ultimately was to create another machine learning system to arbitrate the decisions of Google Translate so that it could produce alternatives, if there was bias that was demonstrated in the output.

[00:12:21]
BH: I think there was another example when Apple launched their payment card with Goldman Sachs. The spending limit was higher for some men than it was for their spouses, for example. I think there are a few high profile cases on that, but Apple turned around and said, 'well, there is no bias in this… we've just let the AI algorithm choose what the spending limit should be.' So it had inadvertently come up with that.

[00:12:50]
JB: Yes. I mean, you've hit the nail on the head there, Bradley, because actually what- people assume that a lack of bias is the application of objectivity to a decision making process. They couldn't be further from the truth. Actually, what a lack of bias is, is the application of subjectively applied societal norms. So it's very - a lack of bias can be very subjectively targeted, and that's what makes it so hard to resolve. The other example is the Microsoft Tay bot that was launched. I don't whether you're aware of that one, but you know, Microsoft launched Tay into, a few years ago, into the bulletin boards and the chat boards, and it was intended to interact with users. Of course, it became racist, misogynistic, far right sympathizer and deeply offensive. It was simply a reflection of the demographic on those chat boards. So it was a reflection back of society, and Microsoft had to pull the plug on that, because it was not acceptable to have a machine behaving in that way. That's what makes it incredibly difficult to manage.

[00:13:58]
BH: So how do you manage that from a legal perspective?

[00:13:59]
JB: Well, there are a number of ways you can manage that from a legal perspective, and I think you need to look at the way in which data are sourced for machine learning solutions, data sets are sourced for machine learning solutions. I guess we're still at the latter end of what I call the Wild West phase in terms of this technology and the position is compounded somewhat by the fact that, you know, we have a lot of very exciting startups that are looking for immediate return on the market and looking for investment. That, I think encourages less than precise decision making in terms of what the sources are that they use - you know, where they get their data from. The first point is that they need to apply appropriate due diligence in relation to that, and they need to make certain that the data that they're sourced, they're sourcing from is suitably representative of the demographic of the problem that they're going to be solving with their machine learning solution. They shouldn't just be licensing in wholesale from third party dataset providers. This is a problem that lots of businesses have found with licensing and facial recognition databases, for example, because they've clearly been compiled as a result of a university study. So all they contain is lots of white caucasian faces. They're not designed for use in the real world. The fact that they're looking for a quick return on their investment means that they take a kind of quick decision to use that and it ends up backfiring because the application that uses the facial recognition isn't as efficient as it would be if it had better data.

[00:15:37]
BH: Yeah, it's when big data needs to be bigger.

[00:15:41]
JB: Exactly right. Yeah, yeah, no, that's right. That kind of leads into another major issue with machine learning, and that's on the intellectual property side and the licensing side. Clearly, you've got to license in datasets that are you know - correctly to make certain they're representative and demographic and not off the demographic and not biased. So there are all those licensing issues, but there are also conceptual issues with intellectual property in that we talked about a trained machine learning system and a machine learning system that isn't trained. The problem with current intellectual property law is that it doesn't allow you to protect the trained element of a machine learning system effectively. You've got copyright law, which covers - and I'm oversimplifying this dramatically here, but you've got copyright law, which covers the kind of authorship and the things that are capable of... typically, it might be pieces of software, and then you've got patents which typically deal with the underlying technology and the hardware and the delta where the value is in these machine learning systems is in the trained element, so you could have a trained system and an untrained system, and for the purposes of copyright, the patent law is exactly the same. They've both got software and they've both got hardware. The only difference is the algorithms are not being trained. So the only way that that can be protected effectively at the moment is through trade secrets. So you've got to keep that information secret. That's a big problem for investors in these types of businesses.

[00:17:18]
BH: What are the main differences between European, UK and US approaches to AI regulation and law? We have listeners from right across the geographies. So can you just summarize what the differences are please?

[00:17:31]
JB: We're still in a process of evolution, Bradley, when it comes to regulation, and the three jurisdictions are taking very different approaches, and I would summarize the – let's start with the US. So the US is very much a mixed bag of- it's a composite of state related laws and there's an evolving federal jurisdiction. The state related laws tend to be focused around the collection of biometric information. So that's information relating to you as a person, as an identifiable person. So there are some quite strict laws in in relation to biometric data collection in Illinois, for example, and a few other states. But on top of that, you have an evolving jurisdiction, which is being managed by the FTC, the Federal Trade Commission, which is, I think, quite admirably taking upon itself to look at AI systems, and it has a broad jurisdiction under the FTC Act to stop misleading or deceptive trade practices within the US.

Very interestingly, quite recently, it took a business called Ever to task – and Ever are the publishers of something called Ever Album, which I think is a- I haven't looked at it, but I think it was a kind of database, an AI solution with facial recognition data in it and Ever made a commitment to its customers that they would delete the data if their facial data, if they asked for it. In fact, they didn't, so they were misleading their consumers and the FTC essentially told them to unwind and untrain their machine learning system and delete the personal data that had been unlawfully used within the algorithm, which is quite a dramatic remedy. In fact, this is a remedy that goes beyond what is available currently in the UK and Europe at the moment. They came, they did come to a settlement, but they were actually forced to unwind their machine learning algorithm as a result of that. So that's what's happening in the United States. European Union has the draft AI Act, which is intended to regulate and prohibit some types of A.I. systems and regulate what they call high risk systems, and these typically tend to be – and I won't go into detail, systems which could cause you damage or injury if they go wrong, so obviously automated cars will be covered by that.

But there are also categories for systems that make decisions about whether or not you're eligible for financial services, or products or insurance or systems that are controlling pollution within the European Union, so that you can get a sense for the types of systems that are regulated. That law piggybacks itself on the European product safety regime and forces providers to actually go through a lot of compliance steps to ensure that they are compliant with the regulations. Obviously that is under review at the moment. What is very interesting about that particular law – although it's admirable that the EU is taking a top-down view of regulation, it's not giving users any rights really substantively to seek redress as a result of breach of those rights, which I think is a relatively peculiar decision on the part of the European Union. Because I think if - you're going to put in place a big law that regulates AI, you want to give people the right to seek redress if they're harmed as a result of that AI.

[00:21:09]
BH: Like they have with GDPR and other data laws.

[00:21:13]
JB: Yes, exactly. You know, we come back to Article 22 of the GDPR, which is the automated decision making part of the GDPR, which I think is actually very interestingly drafted and actually, I think would serve as a really nice core for a European law on A.I. because that looks at the output. It looks at whether or not there have been any legal effects on the user as a result of the AI decision making, or any kind of what the GDPR calls similarly significant effects and gives them a remedy as a result of that if they're harmed. So I count that as a bit of a missed opportunity, and I've got - for the listeners that are interested, we've got- iTech Law is just about to publish a green paper on the European law where we go into the amount of detail about where we think these are, these are problems, but I'm not going to go into detail now on that, but that's my impression on the European position, which is very bold and innovative. But I think misses a trick a bit. Then finally, you've got good old U.K., which is sitting between the US and Europe in a post-Brexit environment. I think the likelihood that the U.K. is going to go down the track of implementing the AIA is minimal. I also sit on the All-Party Parliamentary Group for Artificial Intelligence with members of the House of Lords and House of Commons, as an advisory board expert. Their view is that we're going to take a sector led approach, so we're going to look at different industries. We're going to look at things like financial services. We'll look at competition law, look at intellectual property law and rely on existing laws such as, for example, the Equality Act, which would pick up things like bias in A.I. But we're not going to go down the track, at least at the moment. We're not going to go down the track of having a top down law in relation to the regulation of artificial intelligence.

[00:22:55]
BH: Right, OK well, thank you very much. That’s very comprehensive. It's been great to have you on Tech Reimagined and answer some of the big questions around artificial intelligence from a legal perspective. To all of our listeners, I hope you enjoyed today's episode of Tech Reimagined and thank you for joining. Please show us some love and hit that subscribe button if you like the episode, and don't forget to let your friends and colleagues know about the show. If you have any questions or you want to reach out, please drop us a line at endava.com or use the @endava handle on pretty much any of the social media platforms. We look forward to hearing from you. Until next time.