



# AI LEADERS PODCAST SERIES

## EPISODE #29: POLICY AND REGULATION FRAMEWORKS

### AUDIO TRANSCRIPT

**Roger Taylor** [00:00:00] Is that artificial intelligence is a foundational technology that is going to change the world in a very dramatic way, and that our regulatory structures are not currently up to the task of ensuring that it is used in a beneficial way.

**Ali Shah** [00:00:19] Welcome to the AI Leaders podcast. Hi everyone. My name is Ali. I am Global Principal Director for Responsible AI at Accenture. My role connects me with people who are working to shape the development and use of technology such as AI, and that includes people who are advancing the technology itself. They might be building tools to address bias, discrimination and other problems that AI systems create or exacerbate through to those who are really debating and defining the policies and the regulations that are needed to ensure AI technologies benefit our society in the here and now, and those who are considering what it all means for us, for our children, and for generations to come. Today, my guest is Roger Taylor. Roger is somebody who I have had the privilege of working with for a number of years, and he has been at the center of considering how to regulate AI. Roger, welcome.

**Roger Taylor** [00:01:11] Thank you. Glad to be here.

**Ali Shah** [00:01:13] And Roger, just to help our audience really understand who you are, I really wanted to ask you a question that we all face, which is why are we interested and why are we spending our time and our working life in this area? I really want to understand what your why is. What brought you to this world of policy and regulation? Because that is not really where you started. Your career didn't start there.

**Roger Taylor** [00:01:35] No, not at all. I first really got involved in this area when I was a journalist. I was a reporter for The Financial Times, and I got sent out to Silicon Valley in the nineties to report on what was going on. And that was a real shock, just to understand the scale and the depth of what was happening, the digitization of everything, the work that was going on to build these ever more complex systems. And this day we were just talking about software systems to enable mechanical tasks to be done more efficiently and to manage information. And so, we talked about Oracle and Microsoft and these kinds of companies, and that's where I started. Of course, working for a newspaper in Silicon Valley, you quickly get a sense that the future of your industry is somewhat under threat and so I moved on from that and set up a business with a couple of colleagues.



And this was a data driven business looking at how to use data systems to better understand health care and health care outcomes. And we worked with Imperial College to do this, and it was quite controversial at times. We did a lot of work on, for example, comparing mortality rates between hospitals and what that experience really brought home to me was the power, the power that you have in data. All we were doing is creating narratives, but the power to create that narrative and to decide how the data should be put together, to be interpretable, to be turned into real meaning, is an enormous power. And with that came a lot of questions about, well, you know, why are you allowed to do this? What gives you the authority? What's the ethical basis for you doing this? And we put on quite a lot of work into that, and we set up an ethics committee. We created a code of practice. We established a legally independent group that could hold us to account, and we had many complaints and they, on several occasions, ruled against us and said, No, that's it. You shouldn't do that. Then we changed what we did. And so that was a really good grounding both in terms of the complexity of data and understanding that you can pretty much do anything you want with data if you're unscrupulous, and that there is a need to properly be accountable for what you do. And so, after that experience, I moved into regulation. I worked in regulation in health care, worked with the regulation in education. But I think the thing that, perhaps, the reason I'm here now talking to you is because of my role as being the chair that was asked to set up the Center for Data, Ethics and Innovation. And this was the UK's government response to the challenge that governments around the world have realized,

which is that artificial intelligence is a foundational technology that is going to change the world in a very dramatic way, and that our regulatory structures are not currently up to the task of ensuring that it is used in a beneficial way.

**Ali Shah** [00:04:31] So let's probe your career journey and those changes from journalist to setting up your own business and then leading the establishment of the CDI, the Center for Data, Ethics and Innovation in the U.K. Let's just explore that. I think, from hearing your journey and recognizing the massive disruption that was going to come through the use of technology and data and the opportunity that comes from that, you took that leap. Was that leap something that felt scary? Felt exciting? What was it about that change that you were seeing in Silicon Valley that made you think I want to do that here in the U.K. where you're based?

**Roger Taylor** [00:05:09] I have always found it exciting. I've always found the possibility, the potential to do good, to do really amazing things overwhelmingly exciting. And I've always found that that has outweighed the fear of what might go wrong. That's been my personal response to the situation. I think it is just astonishing what people have been able to do, both in terms of the ability to capture astonishingly sort of fine quality data. If you look at the way the Intergovernmental Panel on Climate Change has worked in setting up information gathering systems across the globe, that the level of detail we can get, to their drilling into ice cores to collect data on what pollen levels were thousands of years ago. It is just remarkable what we can do in data collection. And then more recently with artificial intelligence, we have had this revelation of new techniques that allow us to interpret very, very large datasets in ways that we simply would not have been able to do 15, 20 years ago.



**Ali Shah** [00:06:11] I share the excitement that you have, and it constantly surprises me how much more we can do with data and how much more there is still to explore. Everything that you've just described, the wealth of data that's now becoming available and the analysis that allows us to make predictions and decide outcomes and decide our futures is fantastic. But I think in those early days when you were really thinking about the power of data, especially in health care, it also strikes me that you were thinking about the ethical issues and the questions at a time when there was no real regulation mandating you needed to do this around data. I just wanted to understand what made you set up a committee to look at the ethics of your data use. What was driving you then?

**Roger Taylor** [00:06:53] It was the realization that whether you're using data to drive, you know, a car down on road or using data to diagnose someone, or you're using data to try and assess the impact of CO2 on the climate, what is essential is that people trust it and that it is accurate. That's the first point. The second point is, once you have data, there are always ways that you can do things that undermine that, and you might often be in a situation where you can make a lot more money if you do the wrong thing. And so, there's a regulatory problem there. But what struck me in particular was I don't think we had a very clear idea - and I spent quite a lot of time in 2015 writing about this - we didn't really have a very clear idea about how we went about doing this, how we build trusted systems. And a classic example of this was, of course, the situation at University of East Anglia that caused such mayhem around the policy, the move towards getting international agreement about the need to tackle climate change. And this was where the some of the data that had been used. Organizations that wished to question the validity of the conclusions of the panel, had requested access to data.

And the scientists felt, and I think not unreasonably felt, that these people were not honest in their intent, that they had a specific agenda. And as anyone who's worked closely with data knows, it would not be very difficult if you handed overall all the data for somebody to come back and say, well, I've looked the same data, and I come up with a completely different conclusion. I think you've got it all wrong, which is pretty confusing for the world because the world's got to decide what it's actually going to do, yeah? And so, what's very clear to me is that on the one hand, you cannot have trusted systems unless people can be held to account for how they have used that data. On the other hand, you cannot just have a free for all in which everybody in the world shares what they think the data means and finding the right balance to know that we can trust organizations, that they are held to account, because the nature of AI systems, the nature of deeply complex data driven systems, is it is not obvious what they are doing. And we see this time and time again where corporates come up, even at quite simple levels. So, an example would be, you know, in the US the use of data about demand for health care being taken up as an indicator for people's health care needs. And when it was looked at by others, they quickly recognized that, you know, it doesn't really work that way because white Americans were being more demanding of health care for a given level of need than black Americans, and they had built a system that was essentially biased against a whole group of citizens.



Now no one intended that to happen, but it's very difficult to see what is going on with these systems, and the more and more you put data into them, if you build a very complex credit model, you know, you have a risk that you might be picking out things and discriminate against things that you just didn't know, that that's what you're doing. So, the need to be held to account and the need to actually know what you're doing and be in control of it is a problem that we haven't really quite solved yet. I think we're all still working on it, but how you allow legitimate differences of opinion but at the same time, you can build a trusted ecosystem. That's really that's the crux of the problem for me.

**Ali Shah** [00:10:17] So you've given some examples there and you also mentioned regulation, but before we even come on to say, discussing formal regulatory mechanisms and the work that you've been doing, I think one of the things you've really drawn attention to in what you were saying was that for success, for any business in putting forward a data driven or AI driven product and making some assertions off the back of that data, people need to trust the results and that trust needs to have some structure around it because data does not lead you to immutable truths on its own, it's open to interpretation. How can I go about trusting what you are selling me or telling me? And part of the approach that you were trying to put in place was to try and build that framework. I think that sort of takes us into your role establishing the CDEI. So, the Center for Data, Ethics and Innovation, I think I want to try and probe with you why it wasn't called the Center for AI Ethics and Innovation. Why Data Ethics?

**Roger Taylor** [00:11:16] Sure. And so, the reason for it being data ethics is because you can't really separate these issues. As I say, once you have a lot of data, you effectively have a lot of power. And we need to start thinking of it, thinking about it from that moment.

It's also, I think, partly because sometimes people get worried about definitions of AI and the category of issues that we are concerned about is really as complex data driven systems. And as, for example, the European Union legislation makes clear, you don't actually have to have used machine learning to be covered by the AI act. It's any sufficiently complex algorithm, and, I think correctly, is what they are concerned about. Because the worry is that, we don't know exactly what that algorithm is doing and that therefore there is a risk to people from it having unintended consequences and the potential for unintended consequences if by their nature they're very hard to predict. I mean, no one would have predicted that Google creating a video recommendation system for YouTube to help put in front of people the video they're most likely to want next would open up the gates for a flood of misinformation from hostile foreign states. It probably didn't appear on the risk register when that was being considered, and so we are dealing with technologies that are very sophisticated, incredibly capable, but also capable of enormous mischief. And jumping around to, you know, why is this happening so fast? I mean, it is important to understand the scale of the capabilities here, because it's not just that we have now been able to create computers that can do human tasks so they can see things, they can read, they can understand language, they can read text, they can translate text. We are even building robots now that can sort of pick the weeds out from fields or pick the fruit of the off the vine. And it's not just those capability. It is the fact that they can then do things with a speed on the scale that that a human could not even approach. So, if we look, for example, at say AI systems for checking legal contracts, the issue is not just that they can achieve the same level of accuracy as a human. The key thing is they can do it thousands and thousands of times quicker than a human, you know?



If you take medical diagnosis, diagnostics at the moment, we're bringing in the ability to read scans quite carefully, and certainly in the UK and I think in a lot of countries you can't just simply rely on a robotic reading of a scan. A human has to look at it as well. But I think in time, if we get comfortable with this, that requirement will probably be dropped simply because the speed with which you could get through very large numbers of patients and the capacity increase by doing that is an overwhelming benefit.

**Ali Shah** [00:14:23] I think these amazing breakthroughs that you're describing are opening up new ways to tackle diseases, new ways for us to navigate our day to day lives, and a huge amount of opportunity. But they're also really testing the way that we can have confidence in these technologies, these systems and what companies and organizations are doing. And so, we have alongside that amazing growth in the use and development and use of AI, also seen a growth in the conversation about regulation and different types of regulation. And so, let's talk about that. What do we mean by regulation?

**Roger Taylor** [00:14:58] So just start at the broadest level. There's sort of a few different things going on. There's moves pretty much across the globe to tackle social media. There's very few territories that have not decided they need to get some sort of control on social media. So, we have the Digital Services Act coming in Europe. We have the online harms regulation in the UK. Singapore, for example, has brought in legislation to try and control misinformation online. Various ideas have been brought to the Senate in the US and none of them got through. But there's, you know, they have a particular relationship to that industry. And of course, in China we've seen some of the implementation of an extremely authoritarian approach to control of these mechanisms and a very different philosophy about what we're trying to achieve.

Alongside that you've then got the question of the growing use of AI across the whole economy, and we're seeing it primarily in areas like logistics, in areas like human resource management, in financial markets, in mining, in retail. There're very few industries that are not being shaken up by the power of these systems. And there is a recognition that there are risks. And we don't know. It's hard to know that. I mean, we're all clear that we're biased. For example, it is a definite risk, and a few areas whereby arise as being a potential risk but in other ways is very hard to judge. I mean, in logistics, there's been a concern that in a really efficient rostering and scheduling systems are unintentionally putting workers health at risk by creating a work pressure environment, in which they simply...there are stories about people who got home from work and just collapsed and that this, you know, in fact, because of an overly efficient AI system determining what they think the worker is capable of, which is, you know, a mistake that one hopes a human would not have made it so. So, there's a recognition that these systems are going to result in things that we're not quite ready for this, and we need to have some sort of regulation. But on that broader question, there are some very different approaches. So, notably we have the European Union proposing a new law with a single overarching framework for the regulation of AI across the whole economy. But that stands in quite marked contrast to the approach, say, in the UK, or where places like the US and Singapore and Australia and others are going, which is thinking much more about recognizing, I think, that a single overarching approach may not be the most efficient way of doing it. And certainly, personally I do have doubts about that, partly because the issues that you're worried about are going to be really, really context dependent.



You know, the problems that will occur in the use of AI in credit markets will not have a huge amount of in common with the problems that will result from its use in logistics or in self-driving vehicles. So, there's that worry and there's also the worry that if you have a sort of regulation that specifically about AI, that you create a sort of an unlevel playing field. Because for a long time, you know, for the foreseeable future, we're likely to see some people offering, say, diagnostic services that don't use AI, and some people will be offering diagnostic services that do use AI. But obviously, we want to hold them to the same standard, you know? Getting a diagnosis right for the patient. It doesn't really matter whether you used AI or didn't use AI. So, we want to have a regulator, a regulatory framework that is focused on getting diagnosis right. And it will obviously have to have some very specific requirements for the use of AI, but it is not necessarily a good idea to have a very different set of regulatory bodies and process. You go through that for one type of diagnostic approach as opposed to another.

**Ali Shah** [00:19:03] So I want to probe that point you're making about the regulatory framework and the role of different regulators in this space. I think you and I have spoken extensively before, in my own background for regulation. It comes into play here. The view that strong regulation is needed, but it has to be practical, and it has to be something that helps individuals who are affected by these systems. Organizations are building these systems, companies that are buying these systems to understand what they need to do. And with that in mind, I think you've painted the picture where we have regulations either existing or being promoted for dealing with certain issues. So, you talked about social media and the online safety bill or the DSA, the Digital Services Act coming through. You've talked about the EU Act in particular around a more AI centric focus.

Is it not the case...or let me test this with you... is it not the case that most of these approaches are trying to get at the same thing and taking a risk-based approach? Try to understand, look, does this system, this technology, create some risks that we need to deal with and here's how we deal with it? Or do you think that these approaches are quite rightly distinct in certain areas? So, regulations around social media should be distinct from regulations for financial products or health care products. How do you see that tapestry of regulations working together? What do you think is the thing that joins them together or what do you think is the thing that is creating some tensions? I'm interested in your view.

**Roger Taylor** [00:20:34] Yeah, it's very interesting and it really demonstrates how around the globe we are all learning how to do this. We're all trying to work it out. So, if you take the social media example, it's quite interesting. If you compare the European legislation against the UK legislation and both the sets of laws are trying to achieve exactly the same thing. They are really very similar. There are some differences about certain particular aspects in certain political areas, but broadly speaking, they're trying to do the same, but they've got very different mechanisms. The UK mechanism is more focused on a duty on the organization not to harm people. The European approach involves much more specific set of requirements that a very large platform operator would be expected to implement. And so, that's, you know, there's a legitimate discussion there about which is the right way to do it. And we will all learn from seeing how it works. If you take the other divide I mentioned, I think this is, to my mind, a bit more problematic. And I think it'll be interesting to see how the European legislation develops over the next couple of years. I think there's little doubt that the European legislation I mean, the one thing it has is it's very concrete, it's very clear, whereas in many other territories, it's perhaps a bit vague what exactly is going on and how regulation is developing.



But I think there is a real issue about the deliverability of the European approach without putting undue burden on industry. Because it is you're being regulated for AI separately from how you're being regulated in your business anyway. So, if you're a financial services company, are you really going to have I mean, you've already got a sort of separate layer of regulation around data protection and then have another layer of regulation that's around the use of AI. There's a worry that we're building up a lot of layers of regulation. And the other worry is that the specific knowledge you need to correctly identify the best way to manage risks within a particular market is not a general scale. My worry is that we will see, I think, the framework is proposed by the EU law... EU law is a very sensible framework. It's hard to sort of, you know, look at it, go well it's pretty. You going to do that. That is a pretty sensible way of going about it. But there is a worry that the way it might fail is in its ability to really get into the conflict between regulators who have a very clear understanding of the specific issues in their market and a more horizontal regulator that may have approaches that are just a bit more kind of standardized. But a lot of that will depend on how the regulations are formulated. There are powers in the regulation to, I think, devolve authority. So, you could have a mechanism whereby these responsibilities are, in fact potentially devolved to other regulators. And my own view and the basis on which the Center for Data, Ethics and Innovation was set up in the UK, was that what we wanted to do was to empower existing regulators because the law hasn't, you know, what we want from credit markets, what we want from good health care, what we want from good human resources management and employment practices. These haven't changed because of the technology. We still want the same things and the foundational principles of the regulators in each of those areas. I'll still summarize what it is we want regulated, what we want managed. So, my own view is we need to empower those regulators to deliver in this new environment with.

**Ali Shah** [00:24:07] Just to challenge you on that, Roger. I think as you said, there's still a way to go. The EU Act is likely to be ratified within the next year or so and then come into force a couple of years later. So, there's still lots of deliberation happening around not only the structure but around how to assess risks and deal with those, but also how to implement it. Is it not the case that the approach that's already happening in the UK, for example, with the DCF, the Digital Regulation Cooperation Forum where four of the largest regulators are coordinating? In full disclosure, my previous role, I was part of that conversation supporting some of the activities of the DCF. But is that sort of model of regulatory cooperation also not likely to happen holistically or by design when the EU starts to think about implementing the EU Act because there are existing bits of EU legislation that also apply? And you already mentioned the DSA, there's other regulations coming through, and I think nation states and the EU will be thinking through how to make this work. Is that not the case? And we may end up in a space where regulation cooperation emerges as the only way to drive this through? Or do you think actually there is a real risk that it won't happen?

**Roger Taylor** [00:25:22] I think what looks like them right now - like very different approaches - I think in a two- or three-years' time, we might just see them sort of converging. We might see territories strengthening central bodies to sort of just keep an eye on things who have not gone down a single regulatory approach and we might see the EU, as it were, thinking about how it implements its approach, which effectively becomes not so very different from its history. I think that's the likely outcome. I agree with that.



**Ali Shah** [00:25:52] So let's take that point where we've come to an agreement and just try and push a little further. Is the convergence going to happen around the right issues and the right things? I'm interested in your views around globally all the important distinctions and differences in how different geographies are approaching this, whether it's because of the nature of the problems that they're facing in those areas, because of cultural norms or because of the way that citizens and consumers are demanding different products with different AI challenges. Just help me understand your views around it. Are we starting to coalesce, hopefully in the next year or two around the right issues, or do you think there's some distinctions we should be very careful about? I'll probe a little further.

**Roger Taylor** [00:26:38] Sure. So, there's some differences. So just to pick out a few examples of where there are differences, I mean, the most obvious and stark difference in terms of geopolitics right now is between China and the West. You can see it very clearly in terms of their approach to social media. But you can also see it in the way that data is used in the Chinese economy. We're seeing much greater integration across different sectors of the economy. We're seeing much more rapid adoption of data within, say, health care systems. So, the linking together of finance, retail and social media in China has gone further than it has in the West. So, there's a difference in approach there. And it's to do with, you know, an inclination towards central control, as it were, as opposed to a west, western sort of fear of central control or, you know, or desire to have it private, limited. Then you've got areas where say if you take a problem like bias, you know, that they don't in most territories in the world will worry about bias. But of course, they won't all worry about the same things.

So, in Europe and in the UK, bias on the basis of sexual orientation would be illegal. In many other parts of the world, it isn't. The China problem is perhaps a very specific one, and companies have just got to decide, do they want to be in China with that, and what are the implications for them of being inside China? The bias issue is a rather different type of issue because any company that's working with significant international presence is likely to find itself operating under different regulatory regimes that have some similarities about the things that they're concerned about, but nonetheless some quite important differences. And so that's going to be an interesting problem organizations are going to have to wrap their heads around.

**Ali Shah** [00:28:43] Let's follow that train of thought. Because I think what I'm hearing from you and I would agree, we are seeing the growth of AI technology really driving existing regulations. Regulators have to think about how to put in place the guardrails and the different checks and balances that they feel are necessary to make sure that consumers are protected, citizens have their rights met. There is a level playing field for businesses. We're also seeing specific regulations coming through, like the EU AI Act and also globally. Different countries like China and others will have very particular ways to try and regulate AI. So that picture is a complex one. It's a fabric, a tapestry of regulations and regulatory frameworks that are coming through, perhaps underpinned by some common principles around bias or dealing with bias or discrimination or trying to make sure that there's accountability and transparency. But you make the point really strongly that exactly how to test for and deal with bias will vary country to country perhaps, or sector to sector or regulation to regulation. How do you navigate such a complex set of different requirements? You know, it feels like complex waters, you know? How do you settle those waters with some confidence?



**Roger Taylor** [00:30:07] Yeah, no, I think that is right. I think for me, the starting point is always to come at this problem from the underlying problem. The reason why this regulation is coming along is a problem for management as well as for the world. If the problem is, how confident are we that we actually know how these systems are operating? We know the consequences of it. We know the distributional effects between different groups of people. And that is just as much a problem in the boardroom as it is in the corridors of Whitehall and Government. And so, I think the first thing to recognize is that at the same time as we're seeing this drive towards regulation, we are seeing a move within corporations too that recognizes that their own governance of these AI systems needs to step up significantly. And that's partly because of the fear of regulatory risk, but also fear of other risks. You could be sued for fear of backlash from the public or from employers, employees who feel that the companies are not acting ethically or just the fear that you end up doing something completely unintentional and which is harmful to your to your job and business. So, the recognition that governance needs to increase within corporations, I think is very widespread. And so that's a very good starting point. And what you need to do, as it were, to govern these things effectively is not so very different from what regulators are expecting you to do. I mean, to take a really obvious, simple example, the first thing that EU regs require you to do is to look at your AI systems and just assess how risky they are. Are they in that they fall under that banned category? They fall under that high risk category? Are they a medium risk category? Were they low risk? And, you know, there's a very large number of corporations currently who wouldn't know that information.

They may not even know how many systems they have, know the way this technology has spread through different markets or your marketing department or your HR department, your logistics, you know, all sorts of areas will have adopted little bits of AI. And so, one thing I think everyone recognizes this is a problem they need to get on top of. So that's a good starting point. I also think that...two points to really make about this. The first is that companies quite reasonably tend to, you know, not put a huge amount of effort into complying with regulations before the regulations are finalized and put on the statute because you quite likely will find you've wasted your time. I think the situation is a little bit different. And the reason I think it's a little bit different is twofold. Firstly, because of the scale of risk attached to this and the complexity of the problem and the novelty of the problem are all the good reasons for getting focused on it right now. I think the second reason is the point that you were just making, which is that we are seeing a very fragmented landscape. And if you're dealing with a very fragmented landscape, what you don't want to have is a very fragmented response where you set up structures in Singapore to deal with Singapore regulation, structure in Europe, you know? You want where possible, to come up with a unified approach to managing this risk. And so, the more that you've built a system that can either manage these risks and act according to the company's own values, the better place you'll be able to do that. And the third point, which is, I think one I would particularly stress at this stage, is that I hope what has come out from this conversation is the fact that we're all learning, all working at how to do this.



You know, you only have to start reading the European regulations where you start saying things like, you know, you're not allowed to manipulate people with subliminal techniques and to be clear but what does that mean? I mean, you know what a billboard post has been doing for the last 200 years? I mean, how are we going to define? There's a lot of areas where it's going to be a lot of work to determine what this all means. And that is where it would be incredibly helpful if industries were able to demonstrate to regulators, look, this is how we do it, this is why, this is how we think we're getting it right in order to influence the regulatory thinking and to make sure it doesn't come up with sort of rules and principles that just don't really make any sense. I mean, another one in the EU, I think is the one that you have to make sure you all your data is accurate. Well, that's clearly an unachievable goal for any organization, so I think that would be enormously helpful. You know, for me, the standout case of this is human resources. I mean, we're seeing AI being used very widely within human resources, I think approaches to ensuring that this is done in a way that is sufficiently accurate, to be justifiable, that it's not biased against particular groups. I think there's huge variability in the way that people approach this. I think there are some approaches are pretty optimistic about how tolerant the courts think of what's allowable. It's an area where the industry, you know, as a whole could come together and start to help, I think, shape which we are seeing in your sector we've had we are seeing this in in the UK, in other parts of the world. But I think there's a huge opportunity here for companies and for industries to step forward and try and say, look, okay, here's a suggestion for how we do this.

**Ali Shah** [00:35:47] Roger for today and just for now, I just want to really thank you. It was great to get your perspective. I learned a great deal and I'm sure our audience members will really benefit from the discussions that we've just had, so thank you very much, Roger.

**Roger Taylor** [00:36:00] Thank you.