

Caitlin McDonald:

Welcome everyone to today's episode. We've got Ollie Buckley who is the Executive Director of the Center for Data Ethics and Innovation, which is an independent UK advisory body launched in 2018 to advise on the governance of AI and data-driven technologies. Immediately before this, Ollie led the Digital Charter team in the UK government's department for Digital Culture, Media and Sports, where he was responsible for policy development on the societal impacts of the Internet and new technology. Ollie has over 10 years experience shaping the UK government's digital agenda in the prime minister's strategy unit, cabinet office, government digital service and DCMS. Before that, he was a strategy consultant working with a variety of private and public sector clients around the world. Thanks so much for joining us today Ollie.

Ollie Buckley:

Thank you for having me.

Caitlin McDonald:

A pleasure. And I wanted to kick off the conversation just to talk a little bit about the CDEI because not all of our clients and listeners are based in the UK. So you were set up in 2018 to connect policymakers, industry, civil society and the public to develop the right governance regime for data-driven technologies. What's the importance of having an interdisciplinary approach to developing these governance models? And are there other actors or sectors which play a role?

Ollie Buckley:

So if we think about our headline mission, it's to advise on how we can maximize the benefits of data-driven technology and AI to society and the economy and that's a huge opportunity. But keeping the big picture in mind means understanding the various elements that go into making a good society, to creating a world in which people want to live. So the mission is not to say how do we go about maximizing profits but not to care about how you get there. Nor is it about saying how do we reduce the risks of these new technologies to zero because in that case you get no innovation and ultimately you don't get to reap the rewards and the benefits. So just as in all the other kind of important aspects of society, it involves trade offs and it involves conversations between different interests to make sure that what you're trying to do is optimize the system for all.

Ollie Buckley:

So when it comes to the applications of these technologies and the development of these technologies, we need to understand the business interests and incentives, the kind of elements that are going to really drive innovation and encourage people to bring their creativity to bear. But we also need to be thinking about how the benefits that are realized are getting distributed and ensure that those different voices are heard. And so when we look at these problems, we very deliberately seek to bring in the widest range of voices possible to make sure that we're understanding these challenges in the round.

Caitlin McDonald:

Yeah. And that all make sense that of course, we live in a very broad brush society which perhaps has different ideas about what does it mean to live in a society where we all want to live in. And so I'm curious about some of the mechanisms that you use to get those voices into play, to understand the broad public and industry sector perspective on where we actually want the ship to go as it were.

Ollie Buckley:

Yes. And the challenge, and particularly it's been a challenge for a new organization like ours, is that that is such a massive question. And the reality of these technologies is that they are being founded and developed across just about every aspect of life that you can imagine. So when we were being set up and looking for bodies or attempts to do these things in the past, it's quite instructive because say, an offsite of the example is the approach that the UK took to embryo research and to the development of IVF, and there was a multi-stakeholder process that is widely regarded as being very successful in that it brought the scientists together with faith leaders, with civil society, with the people that would benefit and they designed a framework that helped to make the trade offs in a way that everyone could accommodate and live with. And it's a great example of a successful process.

Ollie Buckley:

And we thought for a moment minimizing the achievement that that was, it's a relatively defined set of technologies. And the application in the instance is pretty clear. You are dealing with an aspect of medicine, you're dealing with questions around fertility and you can sort of draw a boundary around that. Now, when you look at data and AI and applications that range from self-driving cars to cancer diagnosis to the algorithms that fill your newsfeed or serve up the adverts that you see, you realize that you can't simply approach this question in the round. Because the questions that matter, the risks and the opportunities that you need to account for look totally different in a medical setting where this might be the difference between getting a rapid cancer diagnosis or not versus a setting in which what you're worried about is am I getting the most relevant reality TV shows flag to me as I'm sitting at home and relaxing in the evening.

Ollie Buckley:

So I think that one of our early insights is that we very deliberately didn't go about sort of creating an AI Uber regulator, but rather to realize that you needed to address these questions in context. So if you take the example of the approach that we've taken to our AI barometer, which is that the piece of work we've done looking across the landscape and horizon scanning and sort of identifying areas of opportunity and risk, we took a sectorial approach to that in order to make it a sort of manageable task at least initially. We didn't take all sectors, we took five and for each of those five sectors, so that covered things like financial services and recruitment and health and crime and justice and others, for each of those, we brought together stakeholders and great people that operated in those sectors and said in your environment, what are the opportunities? What are the risks? What are the gaps that you see for data-driven technologies in your area? And that's how you can start to make sense of it.

Ollie Buckley:

Then what you can do and what for us I think is a particularly important part of our mission, is then to step back and say okay, so what are the common threads here. What do we see that might be true in financial services and also in law enforcement. What are the kind of common trade offs that people are grappling with. And then we can start to develop frameworks and methodologies that have a more kind of general relevance and application. But at the end of the day context is key.

Caitlin McDonald:

Yeah. I think that is a phrase that is so useful at any point in any process, but I really like that analogy of saying when you're talking about data-driven technologies, it's the difference between talking specifically about fertility versus biomedical research as a whole versus medicine as an entire discipline

and industry and I think that that's really key here. Is that when you talk about technologies that are making decisions about people's lives, you simply cannot escape the trade-offs of digital ethics in any area of life and really that's one of our kind of key fundamental beliefs at the LEF, is that everyone should have a stake in this conversation because whether you are a data scientist or not, these technologies are influencing and shaping your life and so you as a stakeholder and a participant in this system need to have some kind of a voice in where we're going with this. But to your point, it's so broad that it can get quickly overwhelming in terms of how should we actually move forward.

Ollie Buckley:

Yes. And this ultimately I think reflects a change in the way that we collectively use these technologies, that it's a kind of move into the mainstream. Say where once you might have worked in a large company or government or organization of any kind and IT was a thing and there was a department that worried about it and maybe you used it to help you with your day job but it wasn't how you actually conducted your business. It supported it but it wasn't the means of making things happen in the same way that increasingly it is today. And so when you move from the sides into the mainstream, it means you have to account for all the factors that make up mainstream decision-making, just to sort of take analogy from government.

Ollie Buckley:

It's kind of a recognition that as your moving services online and you're enabling citizens to transact with government in new ways, there's been a kind of increasing understanding of actually how you design those systems, how you kind of make the process of applying for a driving license easy has within it like policy decisions. It's not simply a matter of designing a great website, that's easy to use. It's recognizing that the ways in which it is easy to use, the trade offs you're making in terms of how accessible you're making it to different groups, what functionality do you include and what functionality do you exclude.

Ollie Buckley:

These are more than just technical decisions now, these are policy decisions, they affect how people relate to their government just in the same way that decisions about how you sell things online if you're a business to your customers, it's not just about the technical functionality that you're providing, but it's also like an expression of your brand. It is the defining element of the relationship that you build with your customers. So you need marketing people, you need design people, you need all the different elements of an organization to buy into that, because that is the way that you do business.

Caitlin McDonald:

Yeah. And it fundamentally affects, like you're saying, the assumptions that you make about who's consuming your work and how they get access to that then fundamentally changes in a government context, ultimately you may have a right to something but if you have no way of accessing that right, then that's a really serious problem. And equally in a business context, you may wish to engage with a certain group of consumers but if they don't have the ability to access your services, then that's also a huge problem.

Ollie Buckley:

Yes.

Caitlin McDonald:

And that takes me nicely onto some questions around governance and governance can mean a bunch of different things actually. So Ramon Choudhry, who was on my podcast previously mentioned that you have model governance at the mathematical level, you have organizational governance through ethics boards or operational procedures or expectations or potentially policy positions from a non-commercial perspective and you also have societally imposed governance through regulation and standards. How does the CDI approach this idea of governance and where do you think you particularly play a role in that ecosystem?

Ollie Buckley:

So the first thing to say is of course that Ramon's right and that all those levels of governance, if you like, are important and essential elements of the ecosystem. And again, if you think to analogies in other spaces and I often think that that transport is an interesting analogy, you have at the very top level some societal objectives around what transport should mean. Where you care about at the level of connecting communities, you care about the impacts on economic growth than an opportunity. But you kind of come down a level and you think well, but we don't want to achieve those things at any cost, you have to consider the environmental impacts, you have to consider the safety considerations because you want to make sure that people are able to trust the system, to get on a plane and know that it's very unlikely that anything bad will happen if you get on it.

Ollie Buckley:

And then how do you do that? Then you have to go another level down and then you start to say well, what are the technical considerations and specifications that mean that we can be confident that these planes are safe or what are the emissions targets and levels that mean that we're kind of satisfied that there is an acceptable approach to managing the environmental impacts of these technologies. And the deeper into the weeds you go, you will find a vast array of different specifications and standards that have taken many years to develop, but which when it works, and of course it doesn't always and all systems are imperfect. But when it works, you should have an EK system that funnels from those micro specifications and the standards that ladders up to creating a safer transport environment and a less polluting transport environment that ladders up to these are the principles around human connection and economic growth that we were wanting to achieve in the first place.

Ollie Buckley:

And so that kind of stack, if you like, that kind of skate is sort of true in any aspect of society where these technologies are increasingly playing. So to the extent that our role at the CDI is to be able to take the look, the helicopter view of you like, across the landscape, we care about all these elements. I would say that where we add value and intend to add value is thinking about the glue and the connection. What is it that holds these different pieces together? How do you go from the top level of high level kind of principle and objective down through to the mathematical models and model governance that support those ends? And what does that mean for an organization that is trying to kind of use these technologies but to kind of stay compliant with the regulation, to be able to kind of tell that kind of customers and suppliers why they can be confident that they're working in their interests and that they can trust them.

Ollie Buckley:

So we're developing frameworks, methodologies. An early example of that is the trust matrix that we developed and have published which is starting to serve as a map, if you like, or a guide for

organizations that are approaching these challenges and help them to say well, what are our high level objectives? What are the principles and values that we care about? And what is the path that tracks those things down to the level of the analysts that are working day to day, the data scientists that are building the models? Now, in reality, in order to test the usefulness and robustness of those frameworks that we're developing, we are working with partners to test at those different at those different levels.

Ollie Buckley:

So we're doing projects where we're focused on helping organizations to develop the principles that they might care about but we're also kind of working with teams that are doing the analysis and the data science and so they're saying well, what's useful to you day-to-day. So we are working at all those levels of governance, but not as a kind of end in itself, but as a means of helping us to inform the bigger picture and to kind of say well, if we're going to help produce these frameworks and approaches, let's make sure that they are actually a practical use to people.

Caitlin McDonald:

Yeah. And I think that's really...in fact one of the reasons that we started this series is that we could see that there are hundreds and hundreds of AI ethics frameworks available and principals and high level ideas about what should be the case. And to some extent, there are some now practical guidance around you the individual data scientists can use these principles or use these frameworks or use these technological tools to help you, but that kind of glue or that ladder is the piece that we saw that was missing and that was really what we wanted to help achieve through this series was how do you actually put these principles into practice? What are the actual practical steps that you need to take as a business?

Caitlin McDonald:

And that can range from somebody who's working in legal or HR all the way through to somebody who's working on the crunchy end of the data science. Because there's an ecosystem even within a business around who needs to be involved at what time to make this work. So I'm really glad to hear that the trust framework essentially has done my job for me and that everyone can just go and look at that.

Ollie Buckley:

Well, I like to think it's a start, I like to think it's a start. But to your point, the treat is that there are lots of areas where actually we've got pretty clear established values and approaches where the task is just saying how do we translate these things to a digital context. You're not starting with a blank sheet of paper, you just want to make sure that where you're using data and predictive analytics, that you are doing it in a way that's consistent with the values of your business and the way that you you've operated to date. But there are other areas where it's less clear, where it's a bit harder to see how these things translate. And so inevitably we will need to collectively build the tools to ask the questions and work together on the answer. So it's a dynamic field and the kind of governance of it and the principles that we apply to getting that good governance are developing and evolving all the time and that's right, that's how it should be.

Caitlin McDonald:

Yeah. And I think that takes me nicely on to a question around accountability mechanisms. And I think there's something here around watching a field which is essentially in its early phases as it industrializes governance starts to grow around it. So you can look at past examples from public utilities previously

that have grown up in the form of energy or previously railways and even newspapers, you had a kind of consolidation of power around those industries which were then broken apart because it was seeing that too much power in one particular place was not beneficial. And so ultimately those were treated as public public utilities in most case rather than as individual private entities and I suspect that we might see some of that happening around the basic technologies that run our world nowadays with some things that aren't currently treated in that way.

Caitlin McDonald:

But even now, before we reached a place where you essentially moved to that utility framework, you can still talk about accountability in terms of governance and where that accountability sits and how we develop those kinds of accountability mechanisms. And like you say, context is king but there might be something that the CDI has been working on in terms of accountability. Is there anything you can share with us about that?

Ollie Buckley:

Part of my response to that would be to say that the importance of accountability underpins just about everything that we do. I think that really the issues that might scare people the most, when you're talking about AI and data-driven technology and one extreme you have the people that worry about the robots taking over and you get sort of Terminator. On closer to home you have real and justifiable concerns about how personal data is used, where's my ability to use my data for my own empowerment versus it being used against me to manipulate me. All of these concerns ultimately get resolved by clarity around accountability, that that may be something of a simplification. But really I think that accountability is at the root of most of the things.

Ollie Buckley:

You need to be able to point to a person or people and for them to own the outcome, to be responsible for the outcome and therefore responsible for improving it, responsible for taking the hit if things go wrong. And again, like accountability, like governance operates at lots of different levels. So businesses, most senior executives need to be accountable for how those companies behave overall but within those businesses, the level of different departments and different operations, you have people that need to take responsibility at those different levels.

Ollie Buckley:

Now, the key point is to understand that in order to hold people accountable, you need to understand certain things about the impacts of their decisions and how those decisions are made. And so the problem of the black box that is often talked about as one of the fears, particularly around the most sophisticated machine learning techniques where you're saying well, what if we don't understand how did the machine made the decision that it made, what does that mean for accountability. I think the answer has to be well, if a machine can't explain important things about how it's operating in a context where that explanation matters, then do not use that machine. If you need because the law says you do in some instances to justify how you're meeting that credit decisions that your making, how you are determining the prices that you are charging people for their insurance.

Ollie Buckley:

And you need to be able to account for the impact of those so to demonstrate that you have not been doing that in a way that might discriminate against people on the basis of protected characteristics like

sex or race, then it is incumbent on the people that are developing those technologies to make sure that they can answer the relevant questions. But it can never be an acceptable answer to say oh, well I don't know because the machine did it and I don't understand how it works. Ultimately, only people can be held accountable and we need to make sure that it is always sufficiently clear that if you're going to point a finger anywhere, there's a person that you can ultimately point it at. That's how we like mitigate the risk of Terminator, I think.

Caitlin McDonald:

Yeah. And I think actually there's a really nice recent example just in the past few weeks of Uber being held accountable for discriminating against drivers in disproportionately allocating rides to particular types of drivers. And in fact, the case law essentially came out and said well, you might not know how this decision was made but the effect was discriminatory and therefore the process is discriminatory. So if you're going to stand there and say we don't know how this works, well you better figure it out because you're going to have to fix it because you are accountable for that.

Ollie Buckley:

Exactly. And this points to the opportunity as well, is that what the use of these systems can do is identify behaviors or unfairnesses that you were otherwise unaware of. Most often I would like to think these kind of unfair or discriminatory outcomes have not happened by design, not by deliberate human intent, they are the results of other things, including unconscious biases that reflected in historic data and then have become embedded in the way that the system works. The kind of explanations that you can demand of these systems enable you to spot patterns that you might otherwise have missed and then to create a better and fairer system as a result.

Caitlin McDonald:

Yeah. And that's one thing that's come up several times throughout this series, is it isn't just about risk mitigation, it is also about opportunity and that's one of the strongest ones. It's that you can reduce some of the harms of the past through using some of these systems, even though at first of course, you uncover bias but the bias isn't to do with the system itself it's to do with the society and we have tools now to identify that and hopefully to help us overcome them. As much as I would love to talk about this all day, we do at some point, I have to draw this interview to a close. So I think just final thoughts or key takeaway messages for our listeners.

Ollie Buckley:

One of the most exciting things and heartening things about working in this space is the extent to which people at all levels of organizations from the data scientists through to the executives are waking up to these opportunities and also to the risks and that while on the one hand, for those of us that kind of work in the data ethics space day-to-day, it can feel bewildering that there are a hundred different sets of AI principles being produced every week. That kind of good thing, what it shows is that there is an active conversation happening and I kind of sense that there's a realization that with the power of these technologies comes the responsibility to make sure that they're working well but that in doing that, you need to be paying attention, you need to be making conscious choices.

Ollie Buckley:

So I think it is a fantastic thing actually that there is so much going on. And I hope that organizations like ours, but there are of course many other institutions that are out in space, can help us all to help each

other and that through this work and making those connections, we can start to develop the blueprints that help society take advantage of these opportunities. So I very much encourage people to be having these conversations and if they're not happening yet in their organization to take the opportunity to start them and to go about it in a way that very deliberately brings in different voices from around their organizations, from outside too. Because ultimately this only works if it's a multidisciplinary endeavor.

Caitlin McDonald:

Yeah. And thank you so much for that helpful end to this episode, I think that there absolutely is the opportunity to move forward in a way that benefits all of us. And despite the fact that when you do work in this space, it can be quite overwhelming to see all of the problems all the time. The reality is that the fact that the conversations are happening is a positive step and it is a positive direction. So thank you so much Ollie, we really appreciate your participation today.

Caitlin McDonald:

Thanks for listening to the growing digital ethics and practice podcast. You can find out more about the Leading Edge Forum perspective on digital ethics by searching for the phrase Stemming Sinister Tides, where the very first results should be our position paper Stemming Sinister Tides: sustainable Digital Ethics Through Evolution.