9 March 2021

British Institute of International and Comparative Law

KING'S College LONDON

THE DICKSON POON SCHOOL OF LAW

CONTESTING AI EXPLANATIONS IN THE UK

BIICL-King's Panel 24 February 2021 – event summary

https://www.biicl.org/events/11467/contesting-ai-explanations-in-the-uk

Summary

Legal challenges to automated decisions and other developments involving artificial intelligence (AI), especially by public sector authorities, have become increasingly prominent in the UK over recent years. This was a lively discussion about some of the lessons from this experience, with contributions from legal activists, practitioners and academics.

The panellists offered a wide range of examples to illustrate their points, not just notorious cases like Cambridge Analytica or Uber but also lesser-known ones involving algorithms used to decide things like personal care budgets and Covid-19 redundancies. They discussed AI as an emergent technology being tested and shaped in law now, rather than as a theoretical problem for the future.

There was general agreement that transparency is required to enable relevant legal challenges to be brought. In nearly all the examples given, AI systems were being introduced and used under conditions of confidentiality. The panellists pointed to what Robin Allen QC called 'observability' – having knowledge of the fact that relevant decisions are being made – as a precondition for obtaining information about how they are being made. They observed that knowledge sufficient to bring an action is required before legal mechanisms like discovery become relevant.

The panellists discussed a very wide range of relevant law, tending to the view that data protection law (including the so-called 'right to explanation' in Art.22) is just a small piece of the puzzle. The immediate focus of the conversation was mainly on public administrative law and on anti-discrimination law. But it was clear from comments about regulation more broadly that the panel considered that there are urgent legal questions across a broad array of fields including competition and trade. One of the themes that emerged from the discussion was that private as well as public law is relevant, with people affected by private sector as well as public bodies' decisions but far fewer remedies available in private law.

Panellists' remarks tended to the view that there is some way to go in terms of exploring uses of current law to contest AI explanations. The discussion suggested that the law is mainly focused on explanations as retrospective information about AI decisions. The courts are developing principles and rules to apply for new practices, for example fairness and due process in the work of public authorities which has relied on data-driven models for years. Regulators are beginning to understand and clarify that relevant harms fall within their remits in new ways that require clarification.

However successive panellists' comments developed a strong sense that current law is insufficient, and that significantly higher policy priority needs to be accorded to the issues because the law is struggling to keep up with the rapidity with which multiple profound harms are beginning to emerge. Improved transparency was discussed as a necessary but not sufficient condition of efforts to address those harms. There were hints that resource constraints are not just holding back

individual claimants in individual cases but, in a much larger sense, civil society and regulators in being able to address the proliferation of implications across all domains including health, education, the environment, finance, etc. Panellists were concerned to improve definition of relevant harms in collective terms (not just in terms of effects on individual subjects) and to develop 'upstream' challenges questioning the legitimacy of original decisions at political levels rather than just struggling to address the fallout for ordinary people when it has already become too late.

The BIICL-King's team organising this panel discussion plan to pursue this conversation further at a follow-on workshop in the coming weeks.

Discussion

Following BIICL's initial introduction, Prof. Julia Black opened the event by pointing out that this is a vitally important issue because information technologies have become so central to the way our society works and to our systems of government. The motivation for this event is the key question of the role of law in contesting and ultimately holding to account the use of machine learning in different situations and in different social, economic and political contexts. Julia called the panellists' attention to the questions of sectoral distinctions and cross-sectoral themes, and to the challenges of proposing constructive ways in which the law can develop and how to ensure that legal perspectives are better-recognised in policy debates. So not small issues!

Rosa Curling started the panel contributions by explaining why she set up Foxglove with Cori Crider and others. She said that they were motivated by the clear lack of litigation and challenges in the courts on these important issues, despite an abundance of ethical panels, soft law guidance, policies, internal advisory boards, governance initiatives within government departments etc. She said that there is a great deal of value to applying the public law obligations and standards, that have been built up through history and which are rightly a very cherished part of our legal system, to the kinds of data-driven and computer-assisted decision making that are now happening in government.

Rosa's opening case example was Foxglove's work with the Joint Council for the Welfare of Immigrants to contest an algorithm which had been used by the Home Office to categorise entry visa applications without anyone knowing about it for about 5 years. Her first point was that disclosure and transparency is a vital precondition to contestation, but that these are typically very difficult and unreliable: in this instance, the key was a very short piece in the Financial Times and then a chance encounter with a former employee in the Home Office. Without the good fortune to become aware of the algorithm's existence in this way, the case would clearly not have been brought.

Rosa then explained that the case involved a request for judicial review and conventional public law principles including equality and transparency. Their case referred to another before the Supreme Court years before involving a Roma rights NGO challenging discriminating discrimination against Romany people at Prague airport using the Equality Act. She said that the use of AI in this instance was effectively no different, since anyone from a nation listed as risky was simply put on the red pile and a much closer level of scrutiny applied. The Home Office conceded quite quickly and agreed to change their processes; so in that sense this case how the law can successfully contest AI.

But in a larger sense, the fact that we could only do that because of learning about the problem by chance; there clearly needs to be a much more determined push – not just through litigation but in other ways too – to get better transparency in practice. We all know there is lots of legislation requiring transparency. The reality is that it's not happening.

Ravi Naik then introduced himself and his data rights agency AWO, which like Foxglove acts on cases asserting rights in the face of uses of technology and to empower people in this context. He said that the title of the event today is food for thought because normally data rights cases challenge decisions or outcomes, not explanations; whereas this focus encourages us to think about the overall problem of the 'black box' society that Frank Pasquale has written so eloquently about. As Rosa has said, this raises the question of transparency – how do we know what's going on behind the scenes, under the hood? Clearly this is a big motivation for the data protection regime including the bundle of rights to 'explanation'. But ultimately it the need for transparency that's such a regular feature of so many of our cases. Not just transparency for the sake of it, as an end in itself, but because without transparency it's impossible to understand how a decision was reached, to challenge it and its outcome.

Ravi gave two examples to develop his argument. The first was Professor David Carroll's case against Cambridge Analytica on which AWO represented Professor Carroll. Ravi observed that this is something of a *zeitgeist* case. He described how in that case, which involved profiling people based on their political opinions, the first step necessary was to get disclosure of and access to the data. Once successful of that, they were able to show that there had been profiling and to start to challenge that. But then the company folded – something that in itself shows the limits of the law here.

Ravi's second example was Segalov v Sussex Police, which he used to illustrate his point about the law tending to contest individual decisions and outcomes rather than the really significant original decisions, leaving the law to chase technology's tail so to speak. He pointed out that in this case, as with so many others, the technology is often quite rudimentary (excel, not sophisticate Minority Report-style algorithms) but that the principles remain relevant. This case involved someone denied access to a party conference after being labelled for extremism. It was that experience that revealed this label had been applied and enabled us to challenge the fairness of that decision. So again, without transparency we can't contest decisions.

Ravi concluded by agreeing with Rosa that there is a lot of scope for contesting decisions on the basis of current law and that the main challenge is transparency. But he added that there are still clearly gaps in the law that have become apparent, notably in laws allowing contestation of collective (not just individual) harms and of more political-context decision-making about where data-based decision making is being used and how.

Robin Allen QC then took up the thread, starting by explaining his experience in this field with a colleague at Cloisters – Dee Masters – with whom he set up the AI Law Hub in 2018. He mentioned some of the work they have been doing with European regulators, the Trade Unions Congress and the UK Centre for Data Ethics and Innovation (CDEI). He encouraged the audience to look at materials published on the hub, but focused on one particular example of recent work on the use of redundancy selection programmes in the context of Covid-19. The system in question used two main metrics, firstly sales performance figures but then also online automated HireVue interviews which claim to look at thousands of data points and then produce a score. This is a fascinating example of a complete lack of transparency but merely a justification that is put forward. It is understood of course that the pandemic was an important part of the motivation for using this system, but still it begs the question: who actually took this decision, the machine or a person or some combination? This is a question that may or may not be answered in law, depending on whether the case is settled or fought.

Robin agreed on the basis of this and other cases that the issue of transparency is very important. But he added that the experience of the AI Law Hub suggested two additional specific observations. Firstly he said that they are very firmly of the view that Art.22 of the (now UK) GDPR is inadequate in the current context. He quoted from the European Data Protection Board's commentary clarifying that disclosure of code is not required and only meaningful information about the logic rather than a detailed explanation. What this means in the discrimination context remains to resolved, all we have is well-established principles in European law going back to Danfoss about an obligation of transparency and inferences to be drawn if this is not given. But the ECJ in a case called Meister clarified that you can't actually get discovery to secure transparency, so all the law says is that inferences may be drawn if there isn't full transparency. This seems to be an area of law that will need to be developed quite considerably.

Robin added secondly that there are specific moves around this towards developing transparency principles in discrimination law. Bearing in mind the Uber decision in the last few days, it is worth noting a member of the European Parliament has put forward a draft Directive calling for much greater transparency in European law and of course we will have to see how the proposals in the European Commission's White Paper from April last year develop.

Conor McCarthy continued with an example he has been working on which helps to show the practical issues faced by practitioners working to bring challenges and obtain remedies. The example does not concern data protection but rather export control in relation to surveillance equipment and technology. Since this is an ongoing matter he raised points hypothetically as involving not just equipment and software including AI but also broader forms of technology. The main legal test to be met concerns dual use and the prohibition of exports which can be applied for civilian as well as military purposes and which involve a clear risk of serious violations of human rights or international humanitarian law.

Conor's initial observation was that the transparency challenges in this instance are multi-faceted. He observed at least four problems for clients wanting to explore the legality in such cases. First, decisions are being made by officials on the basis of self-certification. Second, officials have very limited means to look behind what they're being told and question board definitions of capability that may not articulate the various ways it might relate to human rights concerns. Third, the technology is dynamic and adaptive and can potentially be changed by the end user (or of course, given the 'black box' nature of this technology, to self-adapt and change in ways humans cannot grasp). And then fourth – and perhaps the most fundamental problem – is that of commercial secrecy. Given all these dimensions, it is immediately clear that transparency is very difficult indeed and requires a great deal of expertise. Decisions in this context are shrouded in a great deal or secrecy and there is a profound lack of transparency.

Conor then turned to the legal tools available to address such a situation. He considered that the Freedom of Information Act (FOIA) and Environmental Information Regulations might be used, but that in all likelihood these will be ineffective. FOIA is very slow, it only entitles access to recorded information which is only very limited here, and anyway there are a whole range of exceptions including commercial confidentiality. Even in the best case, there is no positive obligation to produce an explanation. So then perhaps a good alternative might be the duty of candour which possibly offers better potential for actions to elicit a public authority to detail the technology and its uses precisely. But there are huge difficulties here too, including the need to establish reasonable probability of illegality. And in each approach there's a need to involve private litigation, which brings in CPR, disclosure etc which is above all very expensive and which in reality is not a viable tool

for claimants that are not very well-resourced. So while in theory relevant tools exist, they are very limited and in practice very difficult to apply.

Conor then turned to the question of whether these issues are sector-specific or general. He said, that from a legal point of view, there is undoubtedly a need to follow the issues into the detailed implications for different fields of law. He gave the example of competition law, observing that transparency in the sense of sharing an algorithm between competitors could clearly constitute impermissible collusion. He highlighted work on this in the US by the Office of Technology Research and Investigation (alongside the DoJ). He said that one problem in his view is the lack of tools for private litigants outside data protection law to obtain transparency. And that in policy there is also a clear problem with regulatory resourcing – regulators need to be well-resourced, with in-house expertise to understand the practical problems involved in different algorithmic approaches in different contexts.

Jack Maxwell then started his contribution by quickly introducing the work of the Public Law Project (PLP) and his role. He said that he agreed with the previous points establishing contestation as a two-tier problem, firstly obtaining an explanation (or some form of disclosure in the context of the general transparency problem) and secondly using that explanation in legal processes (obtaining a just outcome or not, pushing for decisions that test the boundary).

Jack offered two examples to help illustrate the challenges on each tier. He observed that these do not address fully-fledged AI but rather relevant principles (since anyway it is widely observed that, although increasingly common, public sector applications of AI in the UK are still limited). The first example he offered was the Savva case which concerned a local authority's use of a model-based system to convert an individual's need into a personal care budget. In this case the Court of Appeal accepted the argument that fairness required the authority to tell a decision subject how the decision was reached, and placed emphasis on the importance of information sufficient to enable understanding of and challenge to the decision. The judgment goes into some detail on what would constitute adequate transparency, disclosure or explanation, saying that the authority should publish the model online and also explain how an individual decision was reached. This was a relevant application of longstanding administrative law principles of procedural fairness. It shows how established common law can be applied to data-driven decision making. A similar recent example from 2018 was the Ames case on legal aid budgets.

Jack's other example illustrated a problem on the other tier, situations where even though an explanation is obtained it does not help address the underlying harm. This is a long running case many people will be familiar with concerning English language testing in Home Office immigration processes. In 2014 the BBC revealed problems with evidence of cheating in the administration of these tests, especially by a particular provider ETS. Eventually the evidence involved tests over a few years, adding up to about 60,000 in total. The test provider used automatic voice recognition software to determine whether cheating had occurred and concluded that nearly every one was suspicious; the Home Office then used that as ground to take action against large numbers of people including detentions and deportations. In this case there was this same initial hurdle of – how was this decision made, what explanation is provided. And it was very difficult with the Home Office saying this is a private provider, they will not disclose and it's not appropriate even for us to have the algorithm because of commercial confidentiality. But then gradually over a year or two a greater degree of transparency was reached with disclosures from both. But then as it turns out that didn't really help that much because the problem turned out to have involved, at least to some significant degree, the integrity of the data inputs and doubts as to whether these had been held securely and not subjected to interference or carelessness. So it remains unclear not only whether an individual

had chested but whether the decisions were lawful as well. These cases are still making their way through the courts and claimants and lawyers and judges are still trying to unpick whether the decisions were fair and lawful.

Jack concluded by saying that he agreed with previous comments about there being room to develop existing legal principles both to try to obtain explanations (transparency) and then to try to contest those explanations, notably in these cases the principles of procedural fairness and transparency. But the jurisprudence is still quite underdeveloped. And not just in the meantime but also in terms of how far this can go, the language testing example shows that the outcomes are not necessarily just (not to mention expensive and time consuming). So what this suggests to us is that the thinking needs to move upstream in two ways: firstly to ensure systems are designed in a way that makes them more easily explainable; and secondly for more policy attention to where automation is used in the first place and where it may be inappropriate for it to be used.

Prof. Lilian Edwards thanked the chair and prefaced her comments by observing that she naturally addresses a more academic perspective. She said firstly that transparency is overrated, in the sense that there is no point in looking to transparency as a complete remedy. Transparency is a means to an end, and it's a very limited means because it implies subsequent agency under level playing field conditions. This is invariably exactly what we don't have. If you're a lucky claimant and you get a Rosa or a Ravi gunning for you, then the sky's the limit. But that, sadly, is not most people. So Lilian reported that she worries about the perception in policy spheres that transparency is a goal in itself and then it will solve problems that it won't.

Lilian said that she feels scarred by the debate around the (UK) GDPR 'right to explanation' from about 2016, which involved long and heated debates in her field. She referred to her work with Michael Veale on this and observed that the jury is still out because there is no case law, but urged people to consider whether a genuine attempt to implement a right to explanation would end up in the form we now possibly have in Art.22. The right applies only to solely-automated decisions made on certain lawful grounds. It is unclear what a decision is exactly, this seems to exclude certain significant stages of relevant processes. Notably only decisions with legal or similarly significant effects are covered, and there is ongoing doubt about whether advertising is covered for example (though now, post-Cambridge Analytica I think we would all agree that targeted advertising – especially in the political context – does have such an effect). Even then you have to work through a morass of special cases and exceptions. Maybe then you start to get near what might feel like some sort of right to explanation, but then there is the problem of the legal effect of the recitals. And above all she returned to the problem of how far there has to be human involvement before a decision is regarded as not being solely-automated; not merely rubberstamping, sure, but what then? And as others have alluded to already, for example in terms of doubts about what meaningful information is, it's not really Art.22 that refers to but the earlier provisions which empower people (Art.15(h)).

Lilian said that a very recent example where we've seen this is in the Uber decision. These questions were involved, Art.22 was considered as a possible means to find out more about the Uber algorithm, to show why drivers are penalised, not getting jobs or only at certain times of day. Not sure what happened to that[1], but the interesting thing is Art.22, Art.15 is not how they got the remedy they wanted – that came from labour law. So even though data protection has been discussed so much, what we've been talking about today as actively useful for algorithmic justice is public law, labour law, discrimination law etc. And as Conor said, that points to the big gaps in

[1] Later clarified that this Art.22 action appears to be underway in the Netherlands.

private law and a need to bolster the job of relevant regulators. So Lilian said that for her the big point is regulation not just via data protection but in a whole swathe of other areas including insurance, credit/loans and various other areas where justice will suffer if regulators aren't resourced to step in on behalf of consumers.

Lilian then said that the levels of transparency and challenge alone are also insufficient but that it is necessary to develop more systemic approaches about social impacts and even to consider the most extreme of remedies in the form of a ban. She said that she recognised this sound illiberal but it would be responsible to consider a list of techniques that are not appropriate for use under certain conditions. She offered the examples of police facial recognition and employers' uses of emotional facial recognition, which is increasingly used in online recruitment and which effectively involves a decision about whether or not someone is trustworthy by looking at their face. She said this is snake oil, there is no evidence that it's accurate. Another example she offered is behavioural targeted advertising, which can be discussed round and round but which in the end we want everyone to stop doing. And another example: 'e-proctoring' which purports to monitor whether people are cheating in distance learning software.

Lilian concluded by observing that we still have some way to go in uncovering all the harms caused by AI-related applications, even though these apparently remain at relatively early stages. She suggested that people look in places we might not expect to look and decisions which never seem to get challenged using transparency arguments. She offered the examples of online content moderation, which is the most obvious and prevalent instance of automated decision making today, and de-platforming, which is becoming topical (not just or even mainly with Trump but with lots of ordinary people whose livelihood might depend on using a certain site). All this makes the transparency discussions in relation to the EU Digital Services Act very interesting. She asked whether there might be an opportunity to link these lines of thought to drafting illegal or harmful content codes of conduct and obtaining not only more transparency but better systems. She concluded by saying of course ultimately it would be great to minimise the need for contestation: privacy by design, fairness by design, safety by design!

Perry Keller then offered comments to wrap up the panel discussion by offering an overall framework in terms of two big trends, or trajectories, that are clashing. He said that the first trajectory is a very long-running one which we might call participatory transparency and accountability. This trajectory involves participation in democracy and legal avenues for the assertion of transparency grounded on public interest. And it very much involves the rickety bridge from transparency to accountability. He said that tracing the trajectory back into history, we can see it has broadened over time. It began with civil litigation and private legal actions, then the development of judicial review. Then other developments like Freedom of Information law. Now more recently data protection law has been strengthened, especially subject access rights. Linking all these together we can observe that this trajectory is one of pragmatism and innovation: putting things together in new ways to achieve transparency and accountability, in public for public authorities but much harder to do, private actors too. The innovation continues, both from recent work on class and representative actions and as we have discussed today from the efforts of lawyers to push the trajectory forward. Perry suggested two observations on this trajectory: it is constrained by time, resources and coordination difficulties across a very broad range of issues; and historically it is evident that it is subject to severe structural and doctrinal constraints because this kind of public activism using the law involves serious burdens, compliance burdens but also risks in disclosure of legitimately confidential information. So, Perry said, these legal avenues are not wide open, they're structured to limit or balance. Doctrinally, judges have often been quite reluctant to innovate – for

example the Google case currently before the Supreme Court which is considering key questions on representative actions and to what extent the more relaxed approach suggested by the Court of Appeal should be permitted.

Perry said that the second trajectory is much shorter but more powerful and disruptive: the introduction of automated decision-making and AI into governmental and commercial spheres. He said this trajectory is really about governance – until now supposed to be just ethical principles but now increasingly filtering into law. This trajectory is marked by a great deal of paternalism, in that it tends to cut into human freedoms in a casually domineering sort of a way. There has not been a lot of public participation in determining what the appropriate benefits are or what the harms might be. Harms have been defined at some remove from 'social reality', if we can call it that. Now that's always true in some sense, but in this case it's happening very fast and the implications are generally recognised to pretty deep. So, Perry said, what we have is a set of very significant potential harms; but to the extent that duties, rights and remedies have been imagined or are being explained in relation to those harms it has been distant from public debate and participation.

Perry said that in his view the discussion today relates to the interaction of these two trajectories: AI has hit the law of public participation in information; and the law of public participation in information is hitting the governance of AI. And in that he said he thought the discussion had tended to confirm the idea that the first trajectory is faltering. Lawyers and the law continue to innovate with energy and passion but there are clear limits to existing systems. For example, as Conor said, FOI carries no right to demand explanation; and without that unless someone happens to have recorded a relevant general explanation in a public authority it's not really fit for the AI era. Similarly, as Lilian observed, there are severe constraints to data protection law here; anyway, subject access is largely about access to existing personal information. So our long-running institutional processes for developing public participation though the law are faltering, they don't really serve the new deep needs to probe automated decision-making.

Perry said that we still have private litigation and judicial review with tools of discovery. Within those, there are serious problems with commercial confidentiality and trade secrecy as well as the practical issues with discovery that have been mentioned. Here at least we have some powers to require explanation. But here we return to the basic problem: what rights are we seeking to enforce, and what remedy? This tends to return us to the difficulty in the other trajectory about defining harms and remedies. Data protection law has been thrust to forefront, excessively so. We're told that we've been provided with a range of rights, and told what they're for; those are your interests, those are your rights, this is how you can enforce them. But this has not supressed the growing unease and debate about intrusions into privacy and uses of data. This is a societal harm. Excessive focus on the individual is blocking the possibility that we resolve, or reach for, or even begin to discuss, the collective societal issue. He said that in the meantime litigation efforts are unlikely to make much headway because of the force of the impact from the second trajectory.

Perry concluded by saying that legal tools do exist to address this and that his research and that of others examines them. He highlighted in particular the need to consider innovations involving intermediary institutions (not just activist organisations and litigation channels but regulators too). He added that ultimately this is a political challenge of finding the will to recognise and address a profound threat to our participatory democracy.

The panel discussion then turned to answering audience questions, guided by Julia. Rosa said that she agreed with comments suggesting that contestation needs to be followed 'upstream' to earlier in decision processes and into policy. She said that I her view there is no democratic mandate for

some of the things currently happening in government and therefore Foxglove, while intending to continue to demand fair processes in the courts, intends to move in this direction. She offered as an example a new case challenging an NHS 'data deal' with a US company called Palantir, which is clear example of major, structural decisions being taken at higher levels with a complete disregard for explanation or justice (for example, with respect to consultation obligations). She said that Foxglove tend to agree that merely contesting how these technologies are used risks missing the point, and that there is a need to include the original question of whether or how they should be used in the first place.

Ravi added two comments about scope to make further use of data protection law to address private actors on data. He said firstly that cases are now starting to test how to apply public law principles of fairness in private settings. And he added that the data protection regime also specifies that data should only be collected for specific limited purposes and not for others, and that cases on which he is instructed have been brought to contest Google's approach in this regard. But he pointed out that yes of course there are limits, notably in this instance to personal data and to the gap apparent between data subjects and data controllers – with a yawning accountability gap increasingly apparent in between. Different regulators are also confusing the situation by taking different approaches, for example to what constitutes a legally significant effect. Finally, he added that much of this might be addressed through the elaboration of the concept of algorithmic auditing being developed by the Ada Lovelace Institute and others.

Robin commented on the discussion about transparency by suggesting a higher-level point is relevant, about what he called 'observability'. He gave the example of young people coming out of education and posting their CV online and then being linked in unobserved ways that lead to them constantly being rejected. He said that there is much more to be said on this point of observability in relation to transparency. And that education needs to be considered in much more details as a field that will be a critical proving ground for the rules that we currently have. Robin then added that he is personally convinced that new and good and apt regulation is required. He said that this means considering regulatory silos, resourcing and focus (QCA have one person working on this, CMA similar until now, EHRC only now starting work). EU has done some useful work although criticised for not going far enough. CDEI seem to have listened to the point that joined-up work is needed. He said that he thinks we need a better risk-based approach to all bases of AI and that ideally the government should take up this role. It's crucial that the impact assessments system should be properly backed, and extended into the equality regime although not yet mandatory there. Finally he addressed the question about risks of the EU becoming another data gatekeeper alongside US and China, answering yes that in his view that is likely to happen because of sheer market size but that he would prefer the EU do it since they are thinking deeply about this, even if a bit slow. He said that he is much more worried about the direction of international trade rules, although this was not discussed yet, especially the potential for conflicts around IP rules entrenched in treaties.

Conor advanced a couple of closing thoughts based on the discussion. He said that transparency may not even be required, or at least practically important, in some cases where the legal outcome did not necessarily rest on knowing what was decided or how but on the outcome. He said that the discussion had caused him to reflect on the importance of legal tools being able to reach beyond public authorities and engage obligations from private actors as well, although his would probably have to be dealt with carefully in view of sector-specific contextual issues. He added that an approach placing stricter obligations on technology designers to ensure outcomes remain within legal limits, in line with Lilian's remarks, also seemed to him worthwhile.

Jack's closing remarks focused on a question about the practical difficulties pinning down precisely what an algorithm, decision, explanation etc are. He said that he thinks that the way forward is probably regulating with reference to principles rather than definitions, as with current data protection and public administrative law. He referred back to his examples demonstrating the relevance of values deeply embedded in administrative law as addressing this point, showing that concrete statutory definitions are arguably to some degree beside the point. Finally he addressed the question of possible comprehensive regulation in this area, pointing to recent CDEI's recent report on algorithmic bias – notably the chapter on public sector transparency which brings together relevant disclosure standards and tools – and suggesting that this might be a good starting point for conversations about codification.

Lilian's closing remarks addressed a question about data portability rights as a possible alternative to explanation, which referred to this as an avenue being tested for telematics data in the car insurance market. She said this is an interesting idea and worth considering further. But she said that ultimately this sort of thing is still tinkering around the edges because of a very large point not yet clearly addressed, maybe because considered obvious: the data protection regime only applies to personal data. And a lot of these issues that the panel has discussed do not involve personal data to any significant degree. She referred to a project starting now with BIICL about government decision-making in response to the pandemic: when to go into lockdown or out of it, how to impose restrictions. She said this is another example of vital decisions based on data. She said that ultimately she agrees that what the panel has been discussing is ways to develop new forms of democratic participation.

Perry's closing remarks concerned Brexit. He said that one of the effects of leaving the EU on this field is to direct our attention more narrowly on the UK and on what is possible and useful here. However he agreed that, regardless of this, transborder data flows and geopolitics will continue to mean that what is happening in the EU, US and further afield will continue to be vital considerations for UK initiatives.

Julia closed the meeting rapidly by thanking the panel and also the audience for their participation. She highlighted the apparent need to continue this conversation in view of the value of developing some of the ideas raised in this event across a broad range. A follow-on workshop is planned for those purposes.

Participants

This was a well-attended event, with nearly 300 people registered to attend from a broad range of backgrounds including people working in civil society organisations, business, government and academia as well as interested members of the public more generally.

Chair

Prof. Julia Black is currently LSE's Strategic Director of Innovation, working in the Law Department since 1994. Her primary research interest is regulation and she has advised policy makers, consumer bodies and regulators on issues of institutional design and regulatory policy. Julia was elected a Fellow of the British Academy in 2015. She was appointed to the Board of UK Research and Innovation in 2017. In 2020 it was announced that she will become President of the British Academy in July 2021.

Panel profiles

Rosa Curling is an international and UK human rights and public law solicitor. She is an expert in both areas of law, having advised and led litigation on issues including freedom of information, privacy, EU law, immigration and refugee law, welfare benefits, access to education services, , access to health services, and climate change and environmental law. She led Foxglove's legal team in their algorithmic justice cases through 2020.

Ravi Naik, the Law Society's 2018 – 2019 Human Rights Lawyer of the Year, has a pioneering practice at the forefront of data rights and technology. He co-founded the data rights agency, AWO in January 2020 and represents clients in high profile and agenda setting data rights cases, including the leading case against Cambridge Analytica for political profiling. In addition to litigating, Ravi is known for his advisory and research work.

Robin Allen QC specialises in employment, equality, and human rights law and is closely involved in the debate about discriminatory technology. With Dee Masters, he founded the AI Law Consultancy. In 2000 they wrote "Regulating for an Equal AI" for Equinet - the European Network of Equality Bodies. The Consultancy is currently working with the Council of Europe and UK regulators, and has just submitted its report on "The Legal Implications of Artificial Intelligence in the Post-pandemic Workplace" to the TUC.

Conor McCarthy's practice encompasses public law, civil claims and international law. He is regularly instructed in complex disputes, both public and commercial, across chambers' core areas of work. He has particular expertise in the areas of competition law, data protection and in claims involving issues of public or private international law or civil liberties. Prior to coming to the bar, Conor was a fellow of BIICL and taught international law at Cambridge University.

Jack Maxwell is a researcher at the Public Law Project. He is particularly interested in how technology is changing the way government makes decisions about people and the channels by which they can challenge those decisions. Jack has a Bachelor in Civil Law from the University of Oxford, and degrees in law and philosophy from the University of Melbourne in Australia. Before coming to the UK, Jack worked as a government lawyer in Australia.

Prof. Lilian Edwards is a Scottish UK-based academic and frequent speaker on issues of Internet law, intellectual property and artificial intelligence. She is on the Advisory Board of the Open Rights Group and the Foundation for Internet Privacy Research and is the Professor of Law, Innovation and Society at Newcastle Law School at Newcastle University. Lilian has been involved with law and artificial intelligence (AI) since 1985. She is Associate Director, and was co-founder, of the Arts and Humanities Research Council (AHRC) Centre for IP and Technology Law (now SCRIPT).

Perry Keller is Reader (Associate Professor) in Media and Information Law at the Dickson Poon School of Law, King's College London. He is a specialist in information law, including legal issues relating to data privacy and security; access rights to information; and freedom of expression. Perry's current research concerns privacy and data protection law responses to governmental and commercial applications of profiling and predictive analytics to citizen and consumer personal data.

Event organisers

This event was organised in partnership between BIICL (led by Prof. Maniatis) and the Dickson Poon School of Law, King's College London (led by Prof. Keller). King's work on this event was supported by the EPSRC under the Trust in Human-Machine Partnership (THuMP) project EP/R033722/1.

Professor Maniatis joined BIICL as Director in September 2018. He was previously Professor of Intellectual Property Law and Head of the Centre for Commercial Law Studies (CCLS), and is now Honorary Professor of Intellectual Property, at Queen Mary University of London. His expertise and research interests cover innovation and trade, trade marks and unfair competition, the interaction between intellectual property and competition law, as well as intellectual property in China.

Dr Irene Pietropaoli is a Research Fellow in Business and Human Rights and joined BIICL in October 2018. She conducts research on corporate human rights due diligence, other aspects of implementation of the UN Guiding Principles on business and human rights, and legislation on modern slavery in supply chain. Irene is also leading BIICL's work on Artificial Intelligence and its links with business and human rights.

Archie Drake is a Research Associate at the Dickson Poon School of Law and a Visiting Research Associate at the Policy Institute at King's College London. His main research interests are in policy and practice for technology in public administration. His prior research work includes investigation into new data with the London Ambulance Service. Archie's professional experience spans international development as well as central and local government in the UK.

Contacts: i.pietropaoli@BIICL.ORG and archie.drake@kcl.ac.uk.