HOUSE OF COMMONS

Science, Innovation and Technology Committee

Oral evidence: Governance of artificial intelligence (AI), HC 945

Wednesday 24 May 2023

Ordered by the House of Commons to be published on 24 May 2023.

Watch the meeting

Members present: Greg Clark (Chair); Aaron Bell; Dawn Butler; Tracey Crouch; Katherine Fletcher; Rebecca Long Bailey; Stephen Metcalfe; Carol Monaghan; Graham Stringer.

Questions 412-538

Witnesses

I: Lindsey Chiswick, Director of Intelligence, Metropolitan Police and Dr Tony Mansfield, Principal Research Scientist, National Physical Laboratory.

II: Michael Birtwistle, Associate Director, AI and Data Law and Policy, Ada Lovelace Institute and Dr Marion Oswald, Senior Research Associate for Safe and Ethical AI and Associate Professor in Law, The Alan Turing Institute and Northumbria University.

Written evidence from witnesses:

- Dr Marion Oswald

Examination of witnesses

Witnesses: Lindsey Chiswick and Dr Tony Mansfield.

Q412 **Chair:** The Science, Innovation and Technology Committee is continuing our inquiry into the governance of artificial intelligence. This morning we are looking at some applications, in particular in the police service. To help us with that, I am very pleased to welcome Lindsey Chiswick, who is director of intelligence for the Metropolitan police—a post that she has had since September 2018—and Dr Tony Mansfield, who is principal research scientist at the National Physical Laboratory, which is one of the laboratories that is funded by the Government and the public sector but independently operated. Dr Mansfield and his team have engaged in an evaluation of the initial deployment of AI in facial recognition technology for the Met. I am very pleased to welcome you both today.

Perhaps I can start off with you, Ms Chiswick. Would you describe how the combination of artificial intelligence and facial recognition is being used at the moment in the Met?

Lindsey Chiswick: I speak as director of intelligence of the Metropolitan police, but I also work on facial recognition for the National Police Chiefs' Council, so if I talk nationally, that is why.

Currently, facial recognition in the Met is being used in two different ways. First, there is live facial recognition, which is what you saw deployed during the coronation. That is when people walk past cameras and are immediately assessed against a watchlist; if there is a match, potentially they will be spoken to. We also use retrospective facial recognition through the Home Office's provision of the police national database, although the Met is procuring our own system.

Nationally, there is a third use of facial recognition that is probably worth mentioning at this point: what we call officer-initiated facial recognition. That is somewhere between retrospective, which is the historical use, and live. It is where an officer in the street has access to a mobile phone and takes a picture of the suspect, which will immediately be compared with a watchlist of images.

Q413 **Chair:** In each of those cases, is it right to say that there is a common denominator, which is a watchlist? In other words, you have a list of not just people but their faces—images of them—that is checked against. It is not an abstract picture, as it were, derived by artificial intelligence, of someone who is likely to commit crime; it is looking for an uploaded picture of a specific individual. Is that correct?

Lindsey Chiswick: Absolutely right. In all those three use cases, it is comparing against a watchlist.

Q414 **Chair:** Could you tell me how that watchlist is compiled?

Lindsey Chiswick: I will take live facial recognition as an example, because it is different between the three different techniques that I just described. In live facial recognition, the watchlist is bespoke and specific to the deployment; it is generated by looking at the intelligence case—essentially, the reason why we are using the technology in the first place. The watchlist is compiled based on crime types and people likely to be in the area at that time who are wanted, whether for serious offences or on court warrants. After that deployment is complete, the watchlist is deleted.

It is very much a bespoke watchlist for the deployment location at that point in time. For the coronation, for example, it was specific and bespoke to that event, and it contained some fixated individuals who we knew were going to go to the coronation and potentially do harm.

Q415 **Chair:** On the artificial intelligence dimension, first, did you design the technology yourselves? Did you write it yourselves, or is it procured?

Lindsey Chiswick: No, absolutely not. It is procured. I will speak openly. It is from NEC. That information is publicly available on our website. NEC manufacture all different types of facial recognition technology, and that is the algorithm that we are using.

Q416 **Chair:** Can you say a little about what the artificial intelligence is in this?

Lindsey Chiswick: Facial recognition technology involves, essentially, biometric matching. Dr Mansfield will probably be able to explain the AI in much better detail than I can, but the AI element of that is where the algorithm is doing that biometric matching. Essentially, it is looking at the watchlist—looking at each image and taking a set of measurements of the face; that is the biometric bit—and comparing it to the measurements of the face of the image that the cameras then pick up.

Q417 **Chair:** Does it differ in any material respect from the e-passport gates, where you put your passport in and the camera recognises whether you are the person on the passport?

Lindsey Chiswick: I don't think it is NEC technology at e-passport gates, and I don't know enough about what they are using to be able to answer that, but behind it is facial recognition technology.

Q418 **Chair:** Briefly, Dr Mansfield—we will come on to it in more detail—are we talking about the same sort of thing, even though they might be manufactured by different people?

Dr Mansfield: The facial recognition in the e-gates is the same sort of thing, although the systems are is manufactured by different companies and developed by different companies. The AI techniques are used to extract the best features for making face comparisons. It is left to artificial intelligence to develop that, rather than being a predefined model.

Q419 **Chair:** So it is a step beyond the e-gates things, which is looking for fixed parameters; it is more exploratory.

Dr Mansfield: Both the e-gate system and the NEC systems probably have a fair amount of neural network artificial intelligence in their development processes.

Q420 **Chair:** They do too; so some of what we are talking about in the police context might have an application for border security, particularly passport control?

Lindsey Chiswick: Potentially, yes. The other thing I should say about the live facial recognition algorithms we use is that it is a closed network. By that I mean there are computers in a van, to put it really simplistically, and there are cameras on the van and that loop is closed. It is not linked up to other cameras, nor is it retaining the information and learning from every face that goes in front of it. That is not happening.

Q421 **Chair:** Okay, so specifically on that, before I turn to my colleagues, it is not that every CCTV camera across London is, as it were, primed to look for a person on the watchlist.

Lindsey Chiswick: Absolutely not.

Q422 **Chair:** You set up a camera in a place that you suspect that the person might be frequenting, and it is that camera and that camera alone, rather than remote ones being accessed.

Lindsey Chiswick: Absolutely. It is limited to the cameras that are specific to that deployment, based on a specific intelligence case that is relevant to the area we are in. For example, we deployed it recently in Camden. In Camden, we know there have been significant amounts of gang crime and knife-enabled robbery, and there was also, you will recall, in January a shooting in Euston—again, gang-affiliated. That seemed to be a good location—there is a strong intelligence case, so we set up cameras in a closed loop back to the van, with a specific watchlist for that deployment.

Q423 **Chair:** That is how it is done at the moment, but in principle could the network of cameras across the country or across the capital be linked to that watchlist so that it would not be a closed loop in the way you describe it, but a general thing. Is it an operational or a policy decision you have made to restrict it, or is it technically impossible to do that broader thing?

Lindsey Chiswick: Technically, I think it probably is possible, but we are governed by the law. Currently, we are operating under common law, primarily, and then a patchwork of other legislation underneath it, and various codes of practice and other information from regulators and advice that we are following as well. Under that, and learning an awful lot from the Bridges judgment and the Bridges appeal, at the moment there must be a solid use case for why we are deploying the technology. We need to understand the questions brought up in Bridges, which are around where we are deploying it and why and what is the proportionality and necessity of that, and who is on the watchlist, and again why and the proportionality and necessity of that.

Q424 **Chair:** My colleagues will want to go into some detail on that case. The subject of our inquiry—its purpose—is the future governance of AI, so those questions of what might be possible are of equal interest to what is currently being done. Thank you for that. I am going to turn to my colleagues now, starting with Dawn Butler, then Graham Stringer.

Q425 **Dawn Butler:** Thank you, Chair. Thank you both for coming in today. Ms Chiswick, I want to pick up some of the responses to the Chair just now. With regard to the data that you collect, you said it is a closed loop. Big Brother Watch has been doing a lot of work around this, so can we confirm that all the data you collect—say, at the coronation of live facial recognition—has been deleted?

Lindsey Chiswick: We do not collect data on a live facial recognition deployment. Imagine you are in the crowd and you walk past the camera; the camera instantaneously compares the faces it sees to the watchlist. It is literally instantaneous. If you are not on that watchlist, not only is your facial image immediately and forever deleted, but it is pixelated. This is what we call a "privacy by design" function, which means that an officer sitting in a van looking at the succession of people coming past the cameras cannot even see the individual faces; they are all pixelated. The only exception is if we get a hit—a positive match—in which case the match can be retained for 31 days, or longer if it is needed for judicial purposes.

Q426 **Dawn Butler:** So if you get a match. Is it not correct that 81% of people flagged by live facial recognition were in fact innocent?

Lindsey Chiswick: I don't recognise those figures. I know they keep popping up, and maybe that is something we can talk about when we come on to the National Physical Laboratory report and what we found out about the accuracy of the algorithm, but, as an example, at the coronation there were two positive matches. That to me shows the precision of the technology, and they were both—

Q427 **Dawn Butler:** But how many false matches were there?

Lindsey Chiswick: Zero.

Q428 **Dawn Butler:** That is the issue, isn't it? In January 2022, the Met deployed LFR near Oxford Circus, and one person triggered on the alert. They were subsequently stopped, but they weren't the person; they were innocent. They provided ID to show they were not the person on the watchlist, but they were subjected to further scrutiny and biometric tests before they were allowed to go. Does that not imply that you are assuming that people are guilty rather than innocent? Why did that person have to provide further information because the live facial recognition was incorrect?

Lindsey Chiswick: I don't have details of that specific incident with me, unfortunately, but I am very happy to provide further details of it later.

Q429 **Chair:** Dr Mansfield, are you familiar with that particular incident?

Dr Mansfield: I am not familiar with it.

Q430 **Dawn Butler:** Does it worry you that there is an over-reliance on a system that is 81% incorrect, so much so that if somebody is stopped and says, "I'm not that person," the response is like, "Prove it to us again."

Lindsey Chiswick: There are two ways I want to answer that. First, I can look back at the last six deployments we have done, and I can tell you that in those six deployments we have had two, three, four true alerts and zero false alerts. I will need to go back to your 2022 example, but that is an example of just how precise the technology is. It is not flagging alerts all the time, whether false or true.

Q431 **Dawn Butler:** Isn't it a question how you measure it? When you say zero false alerts, that is because you are only counting your positive alerts, so what does that mean, exactly?

Lindsey Chiswick: We count true alerts and false alerts, and we test the system throughout a deployment by having what we call a blue list, which is our own members of staff and officers who are on a separate watchlist. They walk past the system every 40 minutes or so to check that it is working accurately and just not detecting anybody.

I said I would answer in two ways. The second answer to your question is what happens when an alert happens. What does the officer do? The facial recognition, or the technology, takes you only so far. It takes you to the point of a potential identification; that's all. It will show the watchlist picture, and it will show the picture of the individual who is a potential identification on the mobile phone device that the officer carries.

Q432 **Chair:** Of the police officer?

Lindsey Chiswick: Yes. It is shown to the police officer. They will look at that, and that is when their normal officer powers kick in. They will look it and think, "Is that a likely match?" If they don't think it is a match, they don't need to do anything. They are not compelled to make the stop.

Dawn Butler: But that is on paper.

Lindsey Chiswick: If they think it is a match, there is no immediate arrest; they go and talk to that person and have an engagement. That is similar to any day-to-day activity of a police officer, who usually acts on suspicion. At this point, they have a very accurate tool to point them in the right direction, so all the other activity you see subsequent to that is standard officer engagement.

Q433 **Dawn Butler:** Standard officer engagement with regard, for instance, to stop and search disproportionately discriminates against young black men in all cases—the Met being one of the highest, with nothing being found 80% of the time. If you are building a system based on an already discriminatory system, you are making that discrimination worse. Do we have a demographic breakdown of the watchlist that the Met currently hold?

Lindsey Chiswick: No, we don't. That is because we do not go after individuals on the watchlist. We do not say, "Let's take that named person and put them on the watchlist." The watchlist is made up of crime types, so the demographic balance of that is not something that is considered. It is about the crime types and the wanted people in the area at the time. Immediately afterwards, as I was saying, that watchlist is deleted.

Q434 **Dawn Butler:** Can we get a demographic breakdown of the watchlist?

Lindsey Chiswick: We have not done that, because there is no policing purpose to retain and then process that information for that reason.

Q435 **Dawn Butler:** If you have a watchlist that is 90% women, it means that you are looking for women in a crowd. If a woman is flagged up, you would say, "Right, that woman is the person on our watchlist. There is a match there." There is a correlation between the demographic breakdown of the watchlist and what you do with that information, so is it possible to get a breakdown?

Lindsey Chiswick: First of all, we would never put a watchlist together that was 90% any demographic differential.

Q436 **Dawn Butler:** But if you do not measure it, how do you know?

Lindsey Chiswick: It is not how it is done. It is based on crime types—people who are wanted for robbery offences in the area, and people who are wanted on court warrants. There is a breakdown of the list of reasons to be included on the watchlist in our policy, which is available and open to the public. We would not go after a specific demographic. That wouldn't be right, would it?

Q437 **Dawn Butler:** Can you provide the Committee with a breakdown?

Lindsey Chiswick: I cannot, because we delete the watchlist within 24 hours immediately after the deployment. There is not a policing purpose to process the data in that way. Technically—

Q438 **Dawn Butler:** Did you say you delete the watchlist?

Lindsey Chiswick: Yes.

Q439 **Dawn Butler:** So the Met hasn't got a watchlist.

Lindsey Chiswick: The watchlist is bespoke to the individual deployment. For each live facial recognition deployment we do, we create a very bespoke watchlist of crime types that mirror the reason for being in that area in the first place. At the end of the deployment, we delete that watchlist, which was bespoke to the deployment. If we were going to deploy again somewhere else for a different reason, we would create a different watchlist of different crime types.

Q440 **Dawn Butler:** So there is no way, at any time, of getting a breakdown. You have only a temporary watchlist; there is never a permanent watchlist.

Lindsey Chiswick: There is not a permanent watchlist for live facial recognition technology. For every deployment, it is based on an individual intelligence case, which is the reason why it is both proportionate and necessary for us to be operating in that area at the time. That watchlist is then bespoke to the deployment. We use it and then we delete it, because there is no lawful reason for us to retain that data. Technically, we could keep the watchlist, but lawfully, we cannot. We delete it each time, because there is no policing reason to process data purely to find a demographic difference within a watchlist.

Q441 **Dawn Butler:** Do you see any risks at all with live facial recognition, other facial recognition, or AI and its deployment in the Met?

Lindsey Chiswick: I completely understand public concern around live facial recognition technology and AI, and there is lots of debate going on around that at the moment. What I have tried to do with the introduction of facial recognition technology in the Met is to do so in a careful, proportionate way and as transparently as possible. We have published an awful lot of information online about how we operate and exactly some of the things that we were just talking about—the distinctions of who can and cannot be on a watchlist, and all the safeguards that we have put in place in order to protect people.

Q442 **Dawn Butler:** What risks have you identified?

Lindsey Chiswick: Gosh, we can run through the risks and the safeguards that we have put in place at every single stage of the operation. Starting right at the beginning, I talked about necessity and proportionality of why we are going there in the first place. This is not a fishing expedition; we are targeting areas where there is public concern about high levels of crime, whether that is knife-enabled thefts on Oxford Street, where they operated before, or whether it is some of the gang-related violence and knife-enabled robbery going on in Camden. Those are some examples of specific use cases.

Carrying that through from why we are there in the first place, there is then the proportionality of the watchlist, following our policy as to who goes on there and why. Then on the day of deployment we have safeguards in place there, so we will not run the deployment if there is something wrong with the technology. The images that are accepted into watchlists need to be of a certain quality. The cameras need to be operating in a certain way to make sure that they are also as accurate as we can make them. Throughout the deployment, special training is given to officers who are going to deploy on the ground, which pulls out some of your concerns.

Q443 **Dawn Butler:** Does the training include further intrusion if somebody innocent has been identified? My last question is whether it concerns you in any way that the Metropolitan police is currently institutionally racist, homophobic and misogynistic, and how that will affect the use of this technology.

Lindsey Chiswick: Can I deal with institutional racism first? Then I will come back to the further intrusion, because I am not sure that I fully understood that. You are referring to Baroness Casey's report and other incidents. What is in that report is truly shocking. We have let the public and our own people down. I was shocked by some of the findings. At the same time, I see it as an opportunity for reform. We must turn things around, and I hope that the very careful way that we have been approaching the deployment of live facial recognition technology is doing just that.

Some of the work that we have done with the National Physical Laboratory in order to really get to the heart of the bias in that technology and understand how we can use AI in a proportionate, fair and equal way has been hugely important to me. That is why we publish what we can online in order to be as transparent as we can be. Before we deploy, we do an awful lot of engagement. We engage not only with commissioners, civil society and various groups, but with different community groups and different demographic differentials across the piece to understand their feelings about the technology. Before we go to a location to deploy, we do something called a community impact assessment, where we go out and talk to members of the community about how they feel about having live facial recognition deployed there.

Q444 **Dawn Butler:** Is it not the case that when Big Brother Watch held up a sign saying that the police were operating live facial recognition, one person avoided where the police were and then got a £90 fine because they covered their face? It was assumed that they had something to hide, when they just did not want to be on a police camera. We police by consent in this country, so if you do not want the police to take an image of your face, which is like a fingerprint, we police by consent, and that is your choice.

Lindsey Chiswick: We do police by consent. In any live facial recognition deployment—those who want to are very welcome to come out and see one in operation—we put notices up to tell the public that live facial recognition is going on in the area.

Dawn Butler: But by the time they get to the notice, their face has already been in the system, because by the time they get to the notice, the camera has already taken a facial picture of them.

Chair: Briefly.

Dawn Butler: Sorry, Chair.

Chair: No, no—it is an important area. Ms Chiswick, just answer that, please.

Lindsey Chiswick: The placement of the signage is prior to getting on to the vision of the camera. The view of the camera is actually quite short. Again, if you want to come out and have a look, you can see the zone of recognition and how far it goes along the road. We also, depending on the deployment and what we are seeking to do, publish on social media what

we are going to be doing, so it is foreseeable to the public. We do not force people to walk through the cameras. I have been on the scene and many people have turned around and decided not to walk through the cameras. That is absolutely fine, but it is used as part of a wider policing operation, usually. There will be police officers on the ground, and if they make an informed choice to stop an individual, that is a policing decision.

Chair: Thank you. Let us take some more colleagues, starting with Graham Stringer and Rebecca Long Bailey.

Q445 **Graham Stringer:** I have just a couple of questions to follow up, Lindsey. You made reassuring statements about the pixelation and not retaining the images. Can you show that? Is it audited or checked by any external body?

Lindsey Chiswick: I can show it to anybody who wants to come and have a look at it. If you are thinking about who oversees us and how the oversight mechanisms work, again, it is quite a complex patchwork. Directly, we have MOPAC—the Mayor's Office for Policing and Crime—which has a seat at my strategic facial recognition board that I run within the Met. Stepping out a little more broadly from MOPAC, there is the Information Commissioner's Office and the Biometrics and Surveillance Camera Commissioner, and we have just started some engagement with the Equality and Human Rights Commission as well.

Q446 **Graham Stringer:** But does anybody go in and check that what you are saying about the technology is accurate?

Lindsey Chiswick: Yes, that was the basis of our commission to the National Physical Laboratory to do the testing. We requested the NPL to do testing to look at how accurate the technology was, and to understand better levels of bias within the algorithm.

Q447 **Graham Stringer:** We will come on to that in a second. Can you show operational benefits?

Lindsey Chiswick: Absolutely. We have had a number of significant arrests as a result of facial recognition technology, including for conspiracy to supply class A drugs; assault on emergency workers; possession with the intent to supply class A drugs; grievous bodily harm; and being unlawfully at large having escaped from prison.

Those are some of the examples that I have brought here today, but there is more benefit than just the number of arrests that the technology alerts police officers to carry out; there is much wider benefit. The coronation is one example of where deterrence was a benefit. You will have noticed that we publicised quite widely in advance that we were going to be there as part of that deterrence effect.

If I recall my time up in Camden when I went to view one of the facial recognition deployments, there was a wider benefit to the community in that area at the time. Actually, we got quite a lot of very positive reaction from shopkeepers and local people because of the impact it was having on crime in that area. When we use facial recognition technology, usually it

will be part of a much wider policing operation, so there will be other stuff happening in the area as well—both some things you might see and some things you might not see.

Q448 **Graham Stringer:** I realise that the answer to this is difficult—you arrest people all the time—but can you show an increase in arrests and convictions against a baseline?

Lindsey Chiswick: I do not think that an increase in arrests is really what this tool is about. This is a precision-based, community crime-fighting tool. To use the terrible analogy of a needle in a haystack, the technology enables us to pick out a potential match of someone who is wanted, usually for very serious crimes, and have the opportunity to go and speak to that person. The results that I just read out to you are people who would still be at large if we had not used that technology. It is not a tool for mass arrests; it is not a tool that is going to give you huge numbers of arrests. It is a tool that is going to focus very precisely on individuals we are trying to identify—

Q449 **Graham Stringer:** I understand that, and I realise that giving an accurate, detailed answer will be difficult, but it is about arresting those wanted people, charging them and hopefully convicting them—if they are guilty—against what would otherwise have happened. I realise that this is a difficult question to answer, but it is important that you can show that there is an improvement along those lines.

Lindsey Chiswick: In so far as those wanted individuals for pretty bad crimes are no longer on the streets, yes. We are doing more work to look at how else we might be able to evaluate the benefits of the technology, bearing in mind that we are still at the quite early stages of what has been quite a careful journey.

Q450 **Graham Stringer:** Sorry to repeat myself—so that is quantified, and you have those figures now.

Lindsey Chiswick: I have figures in front of me about the number of arrests that have been made. It takes time for them to work through the judicial system, so arrests that have been made this year, which is what I have in front of me, will not yet have come through.

Q451 **Graham Stringer:** And you can give the Committee that.

Lindsey Chiswick: Yes, in due course.

Q452 **Chair:** Do you have some figures on arrests this year?

Lindsey Chiswick: Yes, we have deployed vans six times. We have done three deployments: Camden, Islington and then three vans were out during the coronation, so that is three different deployments. There were four true alerts. There were zero false alerts throughout those six deployments, and there have been two arrests made. The others were correct identifications but it was decided that arrest was not a necessary action in the circumstances.

Q453 **Graham Stringer:** Dr Mansfield, there is concern and worry about bias

in the learning of AI built into the system. Can you reassure the Committee, or not, that there is not bias in the system that the Metropolitan police is using?

Dr Mansfield: We were contracted to perform an evaluation of the facial recognition algorithm being used in live facial recognition and the other applications that Lindsey mentioned. The scope of the study was to look at performance in terms of accuracy and the differences in accuracy for different demographics. In doing that, we were running an evaluation using a set of test subjects recruited specifically for the evaluation, in addition to data provided by the Met police—a large reference dataset or a large watchlist, although I hesitate to call it a watchlist, because it is far larger than it would normally be for a watchlist. It is a large amount of data, which allowed us to look at the differences in performance on a demographic-by-demographic basis.

Taking the findings on performance, we found the system to be very accurate. In some of the applications—the retrospective facial recognition and the officer-initiated facial recognition—the system was delivering perfect performance and we were getting 100% accuracy. For those applications when there were no errors for any demographic, that is not showing any bias at all.

For the live facial recognition, which is a slightly harder task for facial recognition, the non-detection rate meant that about 90% of people were being recognised and it was fairly constant across the different demographics. We did observe a difference in performance between some demographics, but the observations we had in terms of the recognition rate—I have to talk a bit statistically—are not inconsistent with the hypothesis that there was no difference at all in performance. If you run an experiment and there is a certain chance of being falsely identified, or falsely not identified, you could have another hypothesis to say that there is no difference between demographic, and the results that we had—the observations we had—were not inconsistent with the hypothesis that there was no difference.

Having said that, we did observe that there was some difference. As I recall, the Asian females were the best performing—had the best recognition rate of the demographics—and the black females had the lowest. But those were really quite close. You could never show that there is no bias. Inevitably, when you run a test—if you are tossing coins to find out whether they are fair, you will always find that if you do so a lot of times, say, 1,000 times, it is not going to be 500 heads and 500 tails for a biased coin. So there is always going to be a small amount.

We looked at that in terms of the other type of error rate—people being falsely identified. On threshold setting, the facial recognition systems have a controllable parameter, and as you change that controllable parameter, it requires more evidence that it is the same face—a higher comparison score. If that is changed in one direction, that will mean that the true detection rate and the false detection will both decrease. Or you could

have it set so that it is more accepting of comparisons, so that the true detection rate and also the false detection rate will increase.

Q454 **Chair:** Dr Mansfield, we will have to be a bit briefer, because we have a lot of questions to get through. If you speak finally on that, I will turn back to Graham.

Dr Mansfield: We find that, if the system is run at low and easy thresholds, the system starts showing a bias against black males and females combined. There is some evidence for that—if the system is operated at certain thresholds, which I believe are outside those that the Met police has been deploying.

Q455 **Graham Stringer:** Perhaps you could give a very brief answer to this. Is there scope for improving the system?

Dr Mansfield: As a scientist who believes in the improvement of technology, there is scope for improving systems, yes.

Q456 **Chair:** At the risk of prolonging this, what kind of improvements? What is the nature of the improvements?

Dr Mansfield: A lot of things are trade-offs. You could have a more expensive system with better cameras that would be able to make better decisions. Every time there is an iteration of facial recognition technology, it tends to give an improvement in performance. We have technology that, in some circumstances, can already give null errors, which is good. That is looking at the technology on its own.

There are also mitigations that might be put in place. We evaluated the algorithm on its own, but the algorithm actually operates in conjunction with an operator or a police officer who also has the information about the other images, so there is scope for avoiding incorrect recognitions at a human level as well.

Chair: Thank you very much indeed. Rebecca Long Bailey and then Aaron Bell.

Q457 **Rebecca Long Bailey:** Lindsey, you mentioned earlier some of the oversight bodies that govern the use of biometrics and facial recognition and some of the regulation and guidance that govern the area. I am sure you will be aware that there is a piece of legislation working its way through Parliament at the moment called the Data Protection and Digital Information (No.2) Bill. It proposes to abolish the office of the Commissioner for the Retention and Use of Biometric Material. It proposes to repeal both the duty on the Government to publish a surveillance camera code of practice governing the use of public space surveillance systems by police and local authorities, and the requirement for a surveillance camera commissioner to oversee it.

That was so concerning that the biometrics commissioner himself wrote to the Bill Committee, stating: "it is worth noting that police accountability in their use of new technology such as facial recognition, voice pattern analysis and other AI-driven capabilities is one of the most

contentious aspects of biometric surveillance yet remains unaddressed, either in the Bill (the focus of which remains solely the regulation of DNA and fingerprints in certain, limited circumstances) or at all." Are you concerned?

Lindsey Chiswick: Taking a step back, what I am concerned about is that this technology is really innovative, new and quite difficult to get our heads around—for all of us, I think—as is AI, as we move forwards. The technology is cross-cutting when it comes to regulators; it cuts across the biometrics commissioner, now combined, as you know, with the surveillance camera commissioner. It cuts across the ICO, the EHRC and so many different places.

I do not think that necessarily more oversight added on and built up is the best oversight. We run the risk of having siloed oversight—oversight for surveillance, oversight for biometrics, oversight from data—when, actually, it cuts across all of that. Currently, there is guidance out there that also crosses over and overlaps a little bit.

So rather than building additional layers of oversight, at a more superficial level, I think it would be great to have simplified oversight, but with the right questions. That is the key—having the right deep dive into how we are using that technology to ensure we are behaving in the way we should and the way we commit to in policy. As a policing organisation, there is also other oversight, such as HMIC, to bring into it as well.

Q458 **Rebecca Long Bailey:** Do you have concerns about the capability that will be available once those offices have been abolished and surveillance monitoring is essentially centralised, for want of a better phrase?

Lindsey Chiswick: I am not an expert on this Bill at all, but from a policing point of view, my understanding is that the role of the surveillance camera commissioner will be absorbed into the Information Commissioner's office. From my point of view, fewer different bodies of surveillance and a more simplistic approach to get to the point of asking the right questions is probably helpful.

Q459 **Rebecca Long Bailey:** On a separate issue, we were discussing the number of images held on the police national database. Do you have a figure as to how many images are currently on the police national database?

Lindsey Chiswick: I don't, I'm afraid. I would have to write to you with that.

Q460 **Rebecca Long Bailey:** If you were able to come back to us with that figure, it would be really helpful.

Lindsey Chiswick: Yes, absolutely. Just to be clear, when we are using live facial recognition technology, the numbers on our watchlist are around 8,000 or 9,000—I think we had just over 10,000 for the coronation. There is a reasoned understanding as to why those crime types or individuals were on that watchlist. For the experiment for the work that Dr Mansfield was talking about, the watchlists or groups were deliberately increased to

allow that to happen. When you look at PND, it is vast, with much, much higher numbers.

Q461 **Rebecca Long Bailey:** I have a final question. On the ground, in terms of the use of this technology, how are police officers and police personnel trained to deploy it?

Lindsey Chiswick: Recognising the concerns we have talked about today, prior to every deployment, the police officers who are going to be involved in the use of the technology have a bespoke training session. That bespoke training session is led by one of my officers who is an expert in facial recognition technology. They run through the kit and how it works. They run through police powers and remind them of the need to make their own decisions at the point of getting that match, which we talked about. We run through the things to watch out for, such as demographic differentials and how age might affect things. We run through how, if a trans person was to walk past, there is a necessity for privacy around that. Those are some of the things they need to be thinking about on the ground. We run through bias—we talk about that. We talk about the findings of Dr Mansfield's report.

Q462 **Aaron Bell:** Ms Chiswick, going back to what you said earlier about proportionality and the guidance and various criteria you operate under, surely there might be instances where the model of linking up all the cameras would be proportionate—if there was a very dangerous escaped criminal or someone who had just committed a serious crime. Would the Met, or perhaps other agencies, be able to use facial recognition technology on a grand scale across the whole capital, or the whole country, if we were looking for one individual and it was a proportionate thing to do?

Lindsey Chiswick: That is a really difficult question, in terms of the proportionality of linking up all cameras to find one person. I think you would also need to weigh into that the likelihood of actually finding that one person by joining up all the cameras. It comes back to this: what is the purpose of what we are trying to achieve, and what is the necessity and proportionality of that? And every time that has to be the starting point and then you make your decision. As I said earlier, technically, yes, that is possible.

Q463 **Aaron Bell:** Is that not what we do already with number plate recognition? I mean, are those systems not linked across the whole country if we are looking for a particular car?

Lindsey Chiswick: I guess the difference is that you are talking now about a biometric, which brings with it additional sensitivities.

Q464 **Aaron Bell:** I will move on to other uses of AI, because we have spoken a lot about facial recognition and you have been very forthcoming in your answers. What other uses of AI are there in the Met? I am talking about things like handling large volumes of data, perhaps from mobile phones, looking at social media, and what is referred to as predictive policing, such as crime mapping and individual risk assessment. Does

the Met use any AI tool in any of these areas?

Lindsey Chiswick: The Met is at the start of experimenting with some of this, so we are using automation in a number of areas—robotic processing, essentially, which is logic-based. If x, then do y, and if there are any anomalies that pop out of that process, then an individual person will look at it.

A number of years ago, we did some AI testing around crime pattern analysis. We learned a lot from it. We didn't then go ahead with it, and we are now looking at how we can incorporate that learning and new technologies—things have moved on quite a lot—potentially to do some more trials in the future.

Of course, we are aware of large language model advances—ChatGPT, for example. There is clearly a lot of potential there, as long as it's approached in the right, careful and proportionate way. At the moment, we are not using on a day-to-day basis that kind of tooling in our work. I think there is real opportunity for policing—for handling large amounts of data, first and foremost, actually.

So, yes, looking to the future, there are lots of opportunity around AI. I think what we are doing at the moment is largely around robotic processing to save time.

Q465 **Aaron Bell:** So there is not an algorithmic process yet? That is what you are implying. I think Durham constabulary did try to create something like that and had to—well, they decided to stop using it, because of concerns about the ethics of it. Is that something that the Met looked into? Did they come to the same conclusions that Durham constabulary did about the ethical risks of doing this?

Lindsey Chiswick: I can't comment on where Durham constabulary got to, beyond knowing that they did some experimentation in that area. I think it is something that the Met are interested in. If we were to go forward with that, it would involve discussions with the various ethicists and ethics panels. For facial recognition, we spent quite a lot of time with the London Policing Ethics Panel. They came up with five conditions—five recommendations—for the ethical deployment of that technology, and we have committed to follow those five conditions—those five recommendations—in all of our deployments. If we were to go forward with another tool like that, I would imagine a similar level of engagement, both with communities and with the ethics panel.

Q466 **Aaron Bell:** If you did go down this route at all, there is obviously a real question about bias potentially being introduced by the dataset you choose. Do you think that it could at all reduce bias, because you take out some of the human bias that I am sure some of your officers have, whether intentional or unintentional, or do you think, as per what Ms Butler said earlier, that it could actually exacerbate bias if it is being trained on historic data that has all that human bias in it?

Lindsey Chiswick: Yes, it could; I think you're right. And it would be one of the things that we would need to pay really close attention to.

As for your point about removing the human bias, I would say that that is something you could potentially demonstrate with facial recognition technology. Rather than having an officer standing at, say, Waterloo station, trying to remember the briefing they got of the 10 photos at a line-up, and then trying to look across and see who might be a match, and using their instinct, actually in a way the precise crime fighting of the facial recognition technology removes that. So, you can put the argument.

Q467 **Tracey Crouch:** Lindsey, you talked about deployment earlier. We have seen an increase, unfortunately, in football-related disorder. Is this something that is being used at football matches or other sporting events?

Lindsey Chiswick: The Met has not used it in that way yet. It's something we'll certainly keep under review, and I can see the potential for its use around football banning orders, for example.

Q468 **Tracey Crouch:** Am I right in thinking, though, that South Wales police may have used it recently at a rugby match?

Lindsey Chiswick: The most recent deployment of South Wales police, I understand, was at a concert recently, where there was intelligence to suggest there was going to be criminality related to that.

Q469 **Tracey Crouch:** If it were to be used, particularly in football, where I could see that there may well be use for it, who pays for it? Would it come out of the national football policing budget, or would it come out of the sports bodies? Who ends up paying for the technology to be used?

Lindsey Chiswick: I don't know the answer to that. The technology itself is paid for, so that does not cost. It would be the cost of the policing operation on the ground. The use of live facial recognition would form only a very small part of that wider policing operation. Live facial recognition technology would be a part of how the policing operation is usually paid for.

Q470 **Tracey Crouch:** I have two questions for you on the national lead work. Staying with the theme of football, do you have conversations with individual clubs about the technology and how it could be used by them?

Lindsey Chiswick: Not yet, but there is clearly a lot of potential around football, transport hubs and airports. The list of potential places for informed and careful use of the technology is vast and, yes, that is another option. That then opens up opportunity around it potentially being owned and run by football clubs themselves. Indeed, if you look in the private sector at the moment, a number of shops are using it already.

Q471 **Tracey Crouch:** My second question has a national lead. My brief mentions a tool that is being used by Kent police. It says that the evidence-based investigation tool is an "algorithm which produces a

probability score of a crime's solvability, which informs decision-making around resourcing and cases". I understand what all those words mean on an individual basis, but can you explain what it means in practical terms?

Lindsey Chiswick: No, I am afraid that I cannot. I am the national lead for facial recognition. I am afraid that AI and tools such as that sit elsewhere. You would really have to speak to someone from Kent police. I do not know that at all.

Q472 **Tracey Crouch:** Okay, thank you. Dr Mansfield, you answered my questions around the main findings of your recently published research. I am interested in what further research in this area is needed. You talked about improvements, but I wonder whether there is any specific research that you would like to do.

Dr Mansfield: There is potential for research exactly on facial recognition. It may be delivering improvements to the existing capabilities, but, of course, NPL's interests are much broader than just facial recognition. We are working on other areas of artificial intelligence. We are involved in the AI Standards Hub, which is led by the Alan Turing Institute and involves PSI, the National Physical Laboratory and the Office for Artificial Intelligence. There are things that are going on in terms of being able to evaluate, or to certify perhaps, artificial intelligence solutions in a wide variety of areas.

Q473 **Tracey Crouch:** Do you know what an evidence-based investigation tool is?

Dr Mansfield: I would say no. As you say, you can look at the words and assume what it might mean. From the sounds of it, it is saying, "Okay, is there a possibility of our working out which crimes are more solvable, or more amenable than others, but I am not totally sure what that in itself means. I would have difficulty if someone said, "Okay, we have such a system, can NPL evaluate it?" Unless I know what it is trying to do, I cannot devise a process.

Chair: We will go to Carol Monaghan. Dawn Butler has another question, and then we will go to Stephen Metcalfe.

Q474 **Carol Monaghan:** Lindsey, I will start with you. You talked a little while ago about how advanced the technology was and how it was difficult to keep up with the advances. We all realise that, as mere mortals, we are not likely to understand how the technology works, but we do understand the implications of using the technology, which is what we are trying to get our heads around this morning.

You talked about not wanting more oversight and about adding layer upon layer of oversight. I want to ask a little bit about legislation. As my colleague, Rebecca Long Bailey, already said, we have a Bill passing through this place at the moment, which is looking to remove the surveillance camera commissioner. You talked about the duplication of different bodies.

It is my understanding that the Information Commissioner's Office has a code of practice, but that is advisory rather than has oversight. That is quite a difference. Where does that leave you if you now have oversight with something that is advisory as opposed to a body that is looking in great detail at what you are doing with people's data and how you are using it?

Lindsey Chiswick: At the moment, I weave into our policy the advice from both the Information Commissioner's Office and the surveillance camera commissioner and biometrics commissioner and the other bodies. The other thing I haven't mentioned yet that we now have in place for the deployment of live facial recognition is authorised professional practice through the College of Policing. It was quite a big step forward when that came in.

I don't see it as making a huge difference to me. I am not going to suddenly change how we are operating. New advice and guidance and policy comes in regularly, and we will always review it and ensure we have adopted, where we can, what that guidance is pointing to. It is not so much a case of not wanting oversight or not wanting more oversight—what I really want is the right sort of oversight, to really get to the right questions, the searching questions, which will give you the assurance that you and the public will want.

Q475 **Carol Monaghan:** Do you think an advisory code is going to give that scrutiny and those searching questions?

Lindsey Chiswick: The code is one thing. The follow-up and oversight and questions that are asked of the Met—of how they are applying it—is another thing. Both of those things need to come together.

Q476 **Carol Monaghan:** Does the ICO do that just now?

Lindsey Chiswick: Throughout the deployment of facial recognition, we have worked very closely with the ICO and been on the receiving end of some very helpful guidance as we have been going forward.

Q477 **Carol Monaghan:** But you keep talking about guidance and advice. That is different from scrutiny. I am trying to establish whether, if the ICO takes on the role that is carried out just now by the surveillance camera commissioner, there is going to be the same level of scrutiny. Working closely can mean guidance and advice. It doesn't necessarily mean an oversight of what is going on and scrutinising the practices.

Lindsey Chiswick: I don't think I can really answer that. I think that is a question for the ICO, as to how they will change if that Bill goes forward and those recommendations—

Q478 **Carol Monaghan:** But they don't carry out that role at the moment.

Lindsey Chiswick: They carry out that role in relation to our data and our data protection impact assessment and how we handle data in relation to facial recognition. The role of the surveillance camera commissioner is to look at how we address his code of practice. This is where it gets

confusing—where there is overlap. You can see now that the confusion is there.

In my mind, if the Bill goes through, and those changes are made, we will also see a shift in how the ICO operates to incorporate those additional measures. That is how I see it. I don't know if that will happen. Again, it is a question for the Home Office and for the ICO.

Q479 **Carol Monaghan:** Dr Mansfield, is there anything you are aware of? I know legislation is maybe not your strength. When you carried out your investigation and produced the report, were there things you felt were slightly grey in terms of legislation?

Dr Mansfield: Not so much in terms of the exact system we were evaluating, but I am involved in international standards for biometrics, particularly those looking at the evaluation of performance. In our committees, we have a piece of work that is addressing the need to be able to audit biometric systems that are used for surveillance. Part of that is because of European legislation. In European legislation, the use of facial recognition for surveillance is a high-risk application. That requires stricter levels of governance. It would be necessary to have more than a performance assessment or to have someone to look in more detail at how systems are being developed and things like that. There is work going on that I am aware of, in bringing together scientists and the evaluation community in biometrics, so that there is more that can be addressed in terms of surveillance systems.

For example, in the evaluation that we did for Lindsey and the Met police, we were just evaluating the algorithms as a black box. We don't make any assumptions about the models inside. We just assess it empirically by having volunteers going through the system. We know what recognition should be made and what recognition should not be made, because we know the ground truth as to the identities.

What we do not have is an audit of the processes in terms of developing the AI that is being used. We don't see what problems they have been trying to address. We don't know what datasets they have used, if they have been using other face datasets.

In other areas at NPL where we are working in AI, we know that there can be issues with whether the training dataset is sufficiently adequate, whether it comes from multiple sources, and whether it is accurately ground truth. There are a number of issues like that that would be applicable to governance of AI systems, and AI systems using face recognition.

Carol Monaghan: Thank you.

Q480 **Dawn Butler:** Lindsey, in response to Tracey, you said that you have never deployed LFR at a football match. How do you decide where you deploy LFR?

Lindsey Chiswick: It is based on an intelligence case, and is part of a wider policing operation. For example, I talked a little bit about Camden, which was the last set of deployments we did. There were high levels of violent robbery. There were—and still are—ongoing gang issues in the area. There was the shooting at the church in Euston in January this year, as some of you will remember.

That, and a number of other issues, pointed towards it being a good location, as part of a wider policing operation going on at the time, to deploy facial recognition in that area. There was a requirement to do so. The other thing I did not mention was violence against women and girls. There is a night-time economy there, and there are some issues around violence against women and girls related to that night-time economy.

Q481 **Dawn Butler:** And none of that applies to football?

Lindsey Chiswick: No. No, it doesn't. It was different for the coronation. The use case, the primary requirement there, was to protect public safety. If there is intelligence that fixated individuals were going to be at the coronation, and we had that image, that was a good use case for that.

If we are already deploying the technology for that reason, and we have other people wanted on court warrants or for other offences, who the intelligence suggests are going to be in the broad area at the time, we can put those on the watchlist as well. That came up in the Bridges judgment and appeal, and that is what we do.

Q482 **Dawn Butler:** Violence against women and girls increases during and after football matches, so I am surprised that it has never been deployed at a football match. How many times has it been deployed at Notting Hill carnival?

Lindsey Chiswick: Once, and that was in the very early days of us learning to use the technology. It was not a very successful deployment. What we learned from that was the sheer number of faces, the density of the crowd, made it very difficult for the technology to operate. There were too many, it was too packed, there were too many faces. We will take that learning and use it going forward.

Q483 **Tracey Crouch:** How does that apply, then, to Wembley Way, with 60,000-plus? Where does it get to the point that there are too many people for it to be deployed?

Lindsey Chiswick: This comes down to recceing in advance where we are going to deploy, using previous knowledge of other events and looking at the best place for that technology to sit. I don't know Wembley, unfortunately. I don't know Wembley Way, but I am not necessarily saying that is where we would deploy the kit. There are a number of places in and around football where it could be deployed.

As part of the planning for that event, the recces would be important—plus the learning that I have just talked about and other learning that we pick up along the way. After every deployment, we consider what we have learned from it and we use that to go forward, but this is early days. You

are absolutely right on VAWG and absolutely right on football, and it is a really interesting point that I will take forward and think about, but this is something we are going to have to build on as we go forward.

Q484 **Stephen Metcalfe:** Good morning. From your view, this is a really useful tool.

Lindsey Chiswick: Yes.

Q485 **Stephen Metcalfe:** I am slightly puzzled about why Dr Mansfield described it as high risk, when you say it is so very accurate.

Dr Mansfield: I said the European Commissioner declared that it counts as high risk.

Q486 **Stephen Metcalfe:** But do you agree with that?

Dr Mansfield: Not necessarily. You can imagine cases where I would classify it as being high risk. In other cases, I perhaps would not.

Q487 **Stephen Metcalfe:** But if it is accurate, how can it be high risk? If it is an accurate tool that identifies people who are on a watchlist, how is it high risk?

Dr Mansfield: It is not my classification.

Q488 **Stephen Metcalfe:** Fair enough. I suppose what I am trying to get to is: is the system doing anything more, on a significantly larger scale, than a well-trained, particularly observant and particularly experienced officer might be able to do by sitting at a good vantage point within a crowd?

Lindsey Chiswick: An officer would never be able to do what that kit does.

Q489 **Stephen Metcalfe:** At scale.

Lindsey Chiswick: At scale.

Q490 **Stephen Metcalfe:** If you had a superhuman officer, is it doing anything more than that, other than identifying people?

Lindsey Chiswick: No, it's not. But an officer could never remember 10,000 faces and identify them in a crowd.

Q491 **Stephen Metcalfe:** Imagine he could for a moment. It would be extraordinary, but that is all it is doing. It is not identifying people outside a list of people who you have put on a watchlist.

Lindsey Chiswick: Absolutely. It only takes us so far.

Q492 **Stephen Metcalfe:** Why is not everyone on the Met police's watchlist on the watchlist that you deploy when you deploy the technology?

Lindsey Chiswick: We have learned a lot from the Bridges judgment and the Bridges appeal. In the appeal and the review, there were questions asked around two deployments, specifically around the who and the where—the who related to the watchlist.

Our policy, and the authorised professional practice that the College of Policing now have, have taken their learning from what came out of Bridges, because there is not a lot out there to go on. We need to have a reason for having people on a watchlist. It cannot just be everybody.

Q493 **Stephen Metcalfe:** This goes back to my difference between the technology and an officer. The Met maintain a database of people they would like to interview.

Lindsey Chiswick: Yes.

Q494 **Stephen Metcalfe:** Why aren't all those people on the list when the technology is deployed? If an officer could remember those tens of thousands, it would be perfectly acceptable for them to go up to someone and say, "I would like to have a word with you and discuss this further."

Lindsey Chiswick: It comes back to necessity and proportionality. The people on the watchlist for a deployment of live facial recognition technology—there needs to be some likelihood of them being in the area when those cameras are deployed. If I am deploying cameras in Camden—it's a transport hub, there are trains, the tube and neighbouring boroughs, and there is evidence of people travelling across London to go to Camden. If there is someone from, say, Newcastle, I would not have them on that watchlist. I understand that your comments are across London but, still, some thought needs to go into that watchlist to look at the likelihood of people on there being caught.

Q495 **Stephen Metcalfe:** But why? Do you think the accuracy of the technology would start to falter if you had too many people on the list?

Lindsey Chiswick: No, not at all. The technology can deal with it. It is to do with the law and the Bridges judgment.

Stephen Metcalfe: Right, and was that a sticking plaster?

Lindsey Chiswick: Was what a sticking plaster?

Q496 **Stephen Metcalfe:** Were those judgments to placate those who opposed the technology, or was there a real operational reason for this? It goes back to the comparison between the technology and a very well-trained, experienced, super-memory officer. It is not doing anything more than identifying those who are potential suspects.

Lindsey Chiswick: I do not think a judicial review operates in that manner, as either a sticking plaster or as something to placate people. A judicial review examines the facts at the time, which are the facts of two very specific deployments.

Q497 **Stephen Metcalfe:** Okay. It is limited in its use. It has a list that is deleted after 24 hours, but you do not delete the main list—you do not ask officers to forget who they are looking for. Why is it only being deployed in Camden and Islington across London? I think you said that it has been deployed six times.

Lindsey Chiswick: I am talking about this year.

Q498 **Stephen Metcalfe:** But if it is such a good tool, why is it not more readily and regularly used?

Lindsey Chiswick: I have talked about the journey we are on at the moment. It has taken a lot of work to get us here in the right way. Now we have the findings of Dr Mansfield's report, that is really helpful, and we do intend to use it more often in future in the right way, with the right safeguards in place.

Q499 **Stephen Metcalfe:** Would you like to change some of the judgments made previously that restrict how this is used?

Lindsey Chiswick: A judicial review is a separate thing. If we look at the law and what we are allowed to operate under, currently we operate under common law, with a patchwork quilt of different laws like the Human Rights Act, the Equality Act and the Data Protection Act underneath that. Then we have College of Policing authorised professional practice underneath that, which was a brilliant step forward and really helpful for police who want to use the technology.

When it comes to the law and AI, regulating each individual technology is probably not the way forward, because the law takes a long time to make, and by the time the law is made and has gone through, the technology will have changed, and it will be out of date. We see that a little bit with the Investigatory Powers Act, for example. In order for law to keep pace with AI and fast-changing technology, broader is probably better, but I can see a future where the law will need to change to keep pace with a whole range of AI opportunities out there, and I guess that is what you are here to think about.

Q500 **Stephen Metcalfe:** Yes, and my final point is about where the future lies. What I am asking you to take a view on is, are we regulating this technology in a different way because it has the word "technology" in it, rather than because of what it actually does, as opposed to the way we regulate the way an officer may behave? Can you see that changing in the future, and would you like it to? What future do you envisage for the use of AI and this sort of facial recognition technology?

Lindsey Chiswick: Oversight is really important as part of the assurance measures for the public, along with our transparency, our engagement and everything else. It is important to do it in the right way, with the right questions that get to the point of how we are using the tech, and then separate officer behaviour and how they go about their day-to-day policing duties. We have oversight from HMIC, MOPAC and a number of other places. I separate those out in my own mind, because to bring them together would be quite difficult; it gets even more confusing.

In terms of the AI, I would like streamlined, clear advice around oversight and what is expected, asking the right questions and not taking where we are now and just building it up and up. It needs something that is co-ordinated and streamlined.

Q501 **Stephen Metcalfe:** Do you think that the technology can potentially be rolled out across an entire city? If the accuracy is as good as it says, is there any reason to restrict how widely it is used? Does it present any more of a threat than an officer does?

Lindsey Chiswick: I do not believe it presents any more threat—I do not even like that terminology—than an officer.

Stephen Metcalfe: All right—risk. I do not what the right word is.

Lindsey Chiswick: I know what you mean. Technically, it's feasible. The future is going to see the use of facial recognition more and more in the private sphere. It is going to see the use of facial recognition potentially by criminals as they try to identify undercover officers, for example, and engage in other nefarious activity. It feels right to me that the police move with the times, look at technology in a careful way and look at how we can make the best of it. At the same time, we police by consent, and a big part of this is public consent to our use of the technology.

Some polling a number of years ago—we need to renew the polling—by a variety of organisations showed that between, I think, 70% and 85% of the public were supportive of police use of live facial recognition technology. But of course that depends on the question asked. There is some really interesting work that can be done around that, so that we can understand and bring the public with us.

Ultimately, it comes down to a privacy versus security debate, doesn't it? Personally, I am relaxed about my image—my biometrics—being momentarily taken and compared against a watchlist; it's pixelated in the back of the van so that people can't see my face. If that helps to prevent crime and stops a wanted offender being on the streets, I am fine with that, but I understand that not everybody is in that position. When it comes to that privacy and security debate, I can contribute to it from the police world, but it is not for me to lead it.

Q502 **Katherine Fletcher:** So much of the use of technology is about the application by humans of it. I want to return to how the faces that are being looked for get into the system and what operational oversight there is, because we have had some concerning incidents in the past, especially around VAWG. I am wondering whether, if somebody is looking for his ex-girlfriend and happens to be the man in the van, he can put that face in, or whether there is a robust framework to make sure that the faces that the system is searching for are only of people that there is a reasonable evidence base for.

Lindsey Chiswick: As reassurance, when one of these operations is put together, there is an authorisation process. There is a form that is filled out with the intelligence case, with the watchlist categories and with different levels of intrusion signed off at different levels, depending on what it is. So we recognise, for example, children, under-18s, and people with disabilities. We talked briefly about trans people earlier. They are examples of areas of extra sensitivity, where extra thought needs to be given.

Q503 **Katherine Fletcher:** Yes, I'm very conscious of that. What I am probably confident in is that you are very robust in the official list that you're searching for. The question is who is checking that—I don't know—Maureen, who is actually loading up the system to look, is only looking for what you think she is looking for.

Lindsey Chiswick: That process I described to you is signed off by a superintendent. If additional images came in—for example, we were deploying and suddenly there was an alert that there was a wanted offender on the street and we needed to get them quickly into the system—that would also have to go through either the superintendent who is the authorising officer for the operation or the gold commander who is overseeing the operation on that day. It may be one and the same person.

Q504 **Katherine Fletcher:** Okay, so the superintendent is in effect in charge of monitoring the fact that only the official list is being searched—

Lindsey Chiswick: Yes.

Q505 **Katherine Fletcher:** And that requires a certain level of technological capability from that individual.

Lindsey Chiswick: The superintendent themselves won't personally be doing that uploading work; that will be undertaken by an engineer.

Q506 **Katherine Fletcher:** Yes, it's the engineer—

Lindsey Chiswick: And the superintendent has oversight of and responsibility for that.

Katherine Fletcher: Perhaps there will be an opportunity to clarify the checking process afterwards—I am conscious of time. Thank you.

Chair: As is evident, we have lots of questions, and probably some more that we would like to follow up in writing with you both, but thank you very much indeed. You have been very comprehensive and wide ranging in your answers. We are very grateful to Lindsey Chiswick and Dr Tony Mansfield for their evidence this morning.

Examination of witnesses

Witnesses: Michael Birtwistle and Dr Marion Oswald.

Q507 **Chair:** We now turn to our next panel of witnesses—pair of witnesses. I invite to the table Michael Birtwistle, who is the associate director for AI and data law and policy at the Ada Lovelace Institute. Joining Mr Birtwistle at the table is Dr Marion Oswald, of the Alan Turing Institute. Dr Oswald is a lawyer as well as an academic and she researches the interaction between law and digital technology, and the human rights, ethics and use of data within policing and the intelligence agencies. Dr Oswald also chairs the West Midlands Police and Crime Commissioner and West Midlands Police data ethics committee. Thank you both very much for coming, and for being here for the previous session; you heard some of the questions and scrutiny that has been on our minds.

I turn to Mr Birtwistle first. Could you give the Committee your assessment of the question that ran through the previous panel: can algorithms and AI, in this case applied to facial recognition, be safely used in policing?

Michael Birtwistle: Thank you, Chair, and thank you to the Committee for having me. I think there are two broad classes of risk associated with the use of biometrics in policing that we have heard today. The first class is about the differential accuracy of those systems, and the second is around the power of facial recognition and biometric systems to undermine democratic rights and freedoms, such as privacy. We think that without a proper regulatory framework there is not a guarantee that facial recognition systems deployed by police will meet reasonable standards of accuracy or that their use will remain proportionate to the risks presented by them.

In July 2019, this Committee called for an independent review of the options for the use and retention of biometric data. The Ada Lovelace Institute decided to respond to that call, and last year published a report that represented the culmination of a three-year programme of work on biometrics governance. That report included public engagement research on attitudes towards biometric technologies, both via a nationally representative survey and a citizens' biometrics council, as well as an independent legal review by Matthew Ryder KC. The report made some urgent recommendations about the sufficient legal framework that would be needed to govern those technologies, both in the public and private sector.

Q508 **Chair:** On that, we are talking about facial recognition, but there is a specific AI dimension to that. Is there a specific concern about the use of artificial intelligence and machine learning with facial recognition technologies, or are the concerns of the Ada Lovelace Institute about facial recognition without the AI dimension?

Michael Birtwistle: There are two parts to that question. Concerns, as articulated in the "Countermeasures" report last year, centre on the use of biometric data as a uniquely personal data type, which is then—to answer the second part of the question—increasingly now used in combination with artificial intelligence technologies. That technology is used to infer not only people's identities but other characteristics. That could be a protected characteristic, but it also could be an internal state, such as whether you are paying attention, what emotion you are feeling or whether you have violent intent—it can make claims about those sorts of predictions. The concern is the combination of both biometric data and the applications that it is increasingly being used for.

Q509 **Chair:** Does that combination not supply some of the answers to the earlier concerns about facial recognition? We know that human beings are prone to biases; if it is a human being who is looking out for people to question or detain, they—consciously or unconsciously—may be biased in the types of people they stop and search. If you have AI that is providing the filter, and is not subject to those same human biases,

does that not inject a greater rigour into the selection of people to be questioned?

Michael Birtwistle: There is an often-cited ambition for that to be the eventual place that those technologies are capable of reaching. But it is not the case in terms of, typically, the accuracy of those systems. Bias, as identified in two CDEI AI barometer reports, is one of the highest risks of these systems. The risk currently goes the other way, towards entrenchment of potential discrimination by these systems.

That is also not the practice of how they are used. The police use a human in the loop currently, and, even under those circumstances, the Court of Appeal held on the Bridges judgment that the use of facial recognition was not according to the law.

Q510 **Chair:** What do you make of the report of the National Physical Laboratory—we have heard from Dr Mansfield—that looked at the deployment by the Metropolitan Police, some of which was experimental, and concluded that there was no significant evidence of bias in terms of demographic groups?

Michael Birtwistle: While it is laudable for the police forces to have submitted their systems to this kind of scrutiny, first, there is no clear legal requirement for police forces to test or demonstrate the accuracy of facial recognition in this way. This is a snapshot of a single version of a single system, set to a very particular setting. The results cannot be generalised to other systems, versions or variations in deployment conditions—poor lighting or whatever, video compression and lower quality watchlist images, which are all covered in section 9.7 of that report as areas for further research.

Chair: Let me take some questions from colleagues, starting with Aaron Bell and then Rebecca Long Bailey.

Q511 **Aaron Bell:** Dr Oswald, what do you see as the main risks associated with AI in the policing and security context? Are they being mitigated, and if they are now, how could they be mitigated?

Dr Oswald: Thank you for the invitation to speak today. I suppose it might be worth briefly summarising what the AI being used by the police is. I think we can probably narrow it down into three or four categories: first, tools that are really there for recording and organising data; secondly, data analytics to try to create insights, trends and connections between people using things such as natural language processing and network analysis. Thirdly and quite importantly, there is risk prediction, which is where the solvability algorithm would come in—trying to predict future happenings or future risk based on past data on either individuals or situations; that is what the solvability algorithm basically is. Fourthly, there are tools to identify people, as you have been discussing this morning.

There is definitely some potential in those tools and we have certainly seen some of them having this potential in the West Midlands committee.

However, there is actually quite limited use of those tools in policing; you might be surprised at the number of forces still using Excel spreadsheets and not the more sophisticated methods. That is down to data basics, such as the fact that data is not necessarily held in a consistent format, and the update for the police national computer is delayed. It is not now to incorporate the intelligence records in the police national database, so a number of bigger forces have invested their own time and money in separate databases. We heard from Lindsey about their investment in a different system for photographs. All data is not really joined up and it is not necessarily consistent, which in itself creates a risk. Also, we do not have a complete picture of what forces are doing with AI machine learning, because there is no central public record or database. We only know from things such as local knowledge, freedom of information requests, House of Lords reports and that sort of thing.

In terms of the risks, there are really no consistent or agreed ethical, scientific evaluation, transparency or quality standards around the technology to decide whether it is any good for a particular policing purpose. Each police force is left to decide whether to develop, deploy or buy technologies. As we heard in the previous session, police forces are asking for more consistent and overarching guidance around that. We also continue to see many examples of AI, sometimes built by commercial providers, promoted for criminal justice purposes, such as AI weapons detection or detecting violent behaviour on trains, without any independent validation or checking of whether those claims are valid. That is particularly concerning in this context.

Chair: Dr Oswald, could we be a little briefer? You have given written evidence; we take as read the evidence that you have submitted. We're keen to have some oral exchanges. Aaron Bell.

Q512 **Aaron Bell:** You set out some of the risks there. Do you think any police force is in a position to evaluate the tools that they are using? Let's leave facial recognition to one side for a minute, because we have talked a lot about it already and in my view all this other stuff is potentially more concerning. We don't know about its efficacy of it; we have not had the sort of tests that the National Physical Laboratory has done, and basically forces are experimenting. One of the advantages of having lots of forces involved is that you can try things out in different areas, but are they in any position to evaluate whether it is any good?

Dr Oswald: I think some forces are, because they have internal expertise to do that or they have bodies they can go to and independent advisers that they can get that sort of advice from. Other, smaller forces may be in a much more limited position. That is why it is so important to set national guidance and national standards, and for the standards to relate to the coercive powers that the police have. The technology is not there on its own; it is related to stop and search, the arrest and investigatory powers that the police are deploying, so the guidance really needs to very clearly relate to that.

Q513 **Aaron Bell:** Do you think all the forces are aware of the ethical risks

involved? Do you think they are aware of the need to keep the human in the loop, as the phrase goes, when they are trialling these things?

Dr Oswald: There is definitely an awareness of that "human in the loop" concept. What it means, though, is quite difficult.

In facial recognition terms, we heard about how officers are given training about the deployment of facial recognition, but when you are presented with a match it is very difficult to decide what to do with it, unless you really understand the settings of the system, what the uncertainties might be and what the risks are for errors and biases.

It is very significant that in the report that the NPL issued, the figures around ethnic disparities are based on a particular face-matched setting. If you look at a setting of 0.56, you will see that of the 33 people falsely identified, 22 were black. At settings that are lower than the recommended settings, ethnic disparities will occur. It is important for the police to be aware of that so that they can deploy those technologies in an appropriate manner.

Q514 **Aaron Bell:** Mr Birtwistle, could I come back to you on the non-facial recognition stuff? What concerns do you have about the use of AI by police in, say, predictive policing or data analytics and so on? Are there examples where it is working well and at the moment you are comfortable with it, or is it something that you are concerned about in the same way that you are about facial recognition?

Michael Birtwistle: The Ada Lovelace Institute has not published research more broadly on the policing use of technology; it has focused on biometrics. I can speak to the proposals in the AI White Paper, which look to existing regulators and existing regulation to govern those technologies.

If we are thinking about whether the measures are going to be there to bring in those accuracy standards or to provide the scrutiny that we have been talking about, the White Paper does not add anything to that picture, other than to provide existing regulators with a set of principles, some central co-ordination and some other risk functions. However, the House of Lords report last year looking at this identified 30 different bodies responsible for looking at policing use of technology, and said it is very difficult to build up a complete picture. It is not clear how the White Paper proposes to improve governance in such diffusely regulated contexts like policing.

Q515 **Rebecca Long Bailey:** Dr Oswald, can you briefly outline what some police forces are doing to increase transparency around the use of AI and data, like the committee that you chair in the West Midlands?

Dr Oswald: The committee was set up by West Midlands police and the West Midlands PCC in 2019 to sit alongside the data analytics work that the force has been doing. It has its own data lab and its own employed data scientists. It was recognised that that created new risks and challenges. We are a voluntary body made up of independent volunteers from all sorts of different backgrounds—law, ethics, computer science,

medical experts, people who have knowledge of the interests of victims—and we take a rolling approach to looking at projects from the beginning, the idea stage, to operational rollout if that happens. Our role is to give pragmatic advice to the chief constable and the PCC and their representatives around the deployment of technology, asking them to think about how these technologies might be deployed operationally and what consequences might occur. All the minutes and papers are transparent, so it is possible through that process to find out a lot about the sorts of projects that the police are considering and the sort of data that is being used.

Another key initiative is the Government's algorithmic transparency recording standard, which has recently been rolled out. A number of police forces have contributed to the pilot of that and demonstrated how, by engaging, more information can be provided to the public around technology being used in police forces. Some research that I did also demonstrated that the police think it could be a really useful initiative to help them to learn from each other, because at the moment they do not necessarily know what other forces are doing and what has been successful and what has not been. By further engaging in that initiative and it even potentially becoming a mandatory process, it has the potential to improve quality in policing technology.

Q516 **Rebecca Long Bailey:** Do you think it should be a mandatory process across all forces?

Dr Oswald: My personal view is yes. There seems to be no reason why it can't be. There are initiatives that you can put in place to protect sensitive information and tradecraft and ensure there are appropriate exemptions. From the research that I have done, I think it could be a really positive step.

Q517 **Rebecca Long Bailey:** Michael, do you think that the efforts that have just been outlined by Dr Oswald are sufficient to address concerns that have been raised about these tools? If not, what more needs to be done? What would you like to see?

Michael Birtwistle: I think they would complement our proposals around comprehensive biometrics governance. I think there is a question about investment. To scale the safeguards and mechanisms that Marion has described, it would take considerable investment to roll that out on a national scale, but that would be a positive thing. I very much echo the endorsement of the algorithmic transparency recording standard and, as it matures, for there to be a mechanism for it to be made mandatory.

Q518 **Carol Monaghan:** Michael, can I start with you? You have already highlighted the link between AI governance and data protection. What are your thoughts on the new Data Protection and Digital Information (No.2) Bill that we are dealing with at the moment?

Michael Birtwistle: Do you mean specifically with reference to policing or more broadly?

Carol Monaghan: You could start with policing.

Michael Birtwistle: I think that with respect to governance of technology by police and particularly biometrics, the DPDI Bill makes the picture worse. As we have heard, it abolishes the roles of Biometrics Commissioner and Surveillance Camera Commissioner, which are obviously held by one individual at the moment, and the code is problematic. The commissioner—commissioners—has written to the Bill Committee stating that, in his view, not all of the functions are being transferred to the ICO or IPCO. Those are some very important provisions around scrutiny and oversight.

It is notable that, when we look at something like Bridges and the way the law currently works, one of the grounds on which facial recognition failed to be lawful there was the question of whether it was done in accordance with the law. When you drill down into the legal technicalities of it, that is saying, "Is there a sufficient legal frame there?" The surveillance camera code was a core part of the legal framework that was examined, and even with that code in place, it was seen as an insufficient framework to be in accordance with the law. That element seems to be being taken away, and in our view that reduces the confidence with which these technologies can be used in policing in a way that gives peace of mind that they are lawful.

Q519 **Carol Monaghan:** We heard from a previous witness about this layer upon layer of oversight and duplication. Do you see the current system as containing duplication? Do you envisage the Bill improving this for users and keepers of personal information?

Michael Birtwistle: I think the claim that it simplifies the regulatory landscape may be true, in the sense that there will be fewer actors in it, but that does not mean that there will be more regulatory clarity for users of that technology. As I said, the removal of the surveillance camera code is one such thing. Our proposal on comprehensive biometrics regulation would centralise a lot of those functions within a specific regulatory function in the ICO that would have specific responsibility for a lot of the safeguards that Marion has described—things like publishing a register of public sector use, requiring biometric technologies to meet scientifically based standards, and having a role in assessing the proportionality of their use. Having all those things happen in one place would be a simplification and would provide appropriate oversight.

Q520 **Carol Monaghan:** Thank you. Dr Oswald, do you have anything to add to that?

Dr Oswald: I agree that simplification does not necessarily mean clarity. Certainly in the policing sector, we need much more clarity about where the responsibility really lies. We have lots of bodies with fingers in the pie, but not necessarily anyone responsible for the overall cooking of the pie. I think that is required here.

Q521 **Carol Monaghan:** Are we getting the right balance between the protection of personal data and the legitimate use of biometric information?

Dr Oswald: It really depends on the context. That is why the regulatory structure needs to be very focused on how the police use data and the different ways AI can be deployed. Sometimes it is deployed in respect of coercive powers—stop and search and arrest—but sometimes it is deployed at the investigation stage and the intelligence stage, which brings on all sorts of different considerations. A regulator needs to understand that, and needs to be able to set rules around those different processes and stages.

Q522 **Carol Monaghan:** Sorry, Michael, did you have something else to add to that?

Michael Birtwistle: I would add to the points made earlier about sufficient future-proofing and flexibility of the law. The way you achieve that is by having empowered regulatory institutions that are properly resourced. That is a concern both with biometrics and generally with the proposals in the White Paper.

Q523 **Carol Monaghan:** I was going to finish, but we are getting rid of two offices and shoving it all into the ICO. Is this an attempt to save money? What is it really trying to do?

Michael Birtwistle: I cannot comment on the rationale for doing that. I think that the importance and benefits of these technologies could be huge. If we want to secure them, we need public trust; we need those technologies to be trustworthy, and that is worth investing regulatory capability in.

Q524 **Katherine Fletcher:** Thanks for coming, Dr Oswald. The Government have an AI White Paper out at the moment. I just want to pivot to what is in that and what we do for the future. You have highlighted in your submission the need for a "context-specific and nuanced approach". I wonder if you could perhaps define that a bit better for us and cross-reference it to the Government's White Paper. Do the proposals in the White Paper meet what you seek to achieve?

Dr Oswald: Of course, yes. I am definitely in favour of an approach that addresses how AI is being used in the policing sector, which is what I mean by a context-specific approach, and really bringing in the key principles of criminal justice. But I am not in favour of relying only on high-level principles and hoping that everyone will be able to interpret those and what they mean, particularly in really high-stakes environments such as policing.

The White Paper is a good first step, but we will need to interpret those principles and find out what transparency really means for policing and AI. What does explainability really mean? What does a police officer need to know when they are using a facial recognition tool, for example? What do they need to know when they are using a solvability algorithm? Those things are not necessarily the same for each tool. Without that, we will potentially continue to see police forces deploying and rolling out different tools in a very individualised way without that consistency of approach.

Q525 **Katherine Fletcher:** Could you give us a specific, for instance, of the next level of detail under the White Paper that you would like to see in this context?

Dr Oswald: If we take, say, a tool that is trying to build connections between individuals to identify new members of an organised crime gang, for example, you might build that data based on intelligence information, or things like non-crime data.

Q526 **Katherine Fletcher:** So that might mean them all being at the same bus stop or something.

Dr Oswald: Yes, exactly, or hanging around on the same street, being seen in the same vicinity or having a link via a phone number or that type of thing. The issue with intelligence is that it does not necessarily mean something bad or something good; it is just a piece of information that you have to interpret. If you then try to put it into a data analytics tool, what happens is that it tends to become either bad or good, and the nuance of it disappears.

You could end up with quite a misleading output that says just because somebody is on this street, they are therefore bad and linked with these other bad people. It is that sort of detail that really needs to be thought about before these tools are built, so you understand the data input and what it really is, and then you build the tool appropriately so that when you get the output, you really know how to use it and you do not necessarily rely on it inappropriately.

Q527 **Katherine Fletcher:** Can I play that back and make sure I understand you—the principles in the White Paper you are comfortable with? It is the practical translation of those principles into a number of fields that is the area of direct concern.

Dr Oswald: Exactly. I think that the expert knowledge that you need to do that is really important here.

Q528 **Katherine Fletcher:** As someone who started my career helping people to get off leatherbound ledgers and into computers, it is that classic transition of—shall we say?—poo in and poo out, perhaps. Thank you for that.

Michael, I will switch to you. Just to understand where we are with the Government's AI White Paper, you are calling for a clear governance framework. That would appear to be a more broad-reaching thing than is within the White Paper at the moment. Will it help to set the scene for such a framework, or will it actually deliver the framework? What is your take on the White Paper?

Michael Birtwistle: To be clear, our call for biometrics governance is probably complementary to what is in the White Paper. The White Paper looks at existing regulation; we are saying that that governance does not exist for biometrics currently, and it will not be added by the White Paper, so it needs to happen separately.

Looking at the White Paper more broadly, there are a lot of advantages to the context-specific approach that it sets out. The real tests for it will be around implementation and—if you will forgive the alliteration—we have three lenses or tests that we will be looking at it through and talking about publicly soon.

First is about coverage. Some domains do not have a domain-specific regulator, like employment and recruitment, so it is not clear who will be responsible for making sure that use cases in that sector accord with the principles. As I said, that is also a problem for diffusely regulated sectors: it is not just policing; education might be another where we have different regulators looking at different bits of the system, but not necessarily with overall oversight.

The second test is about capability—institutional capability. Do regulators have the resources and powers that they need to deliver on the principles? There are some significant questions there about the amount of resourcing needed to govern appropriately a general-purpose technology like AI.

The third one is criticality. Will we have enough action on urgent risks? That is where the question about general purpose AI like ChatGPT comes in, and whether we need a faster response to that ahead of the implementation window of over a year that the Government set out.

Q529 **Katherine Fletcher:** Understood. The basic principles are one thing, but is anything missing, apart from more detail on the implementation perhaps? Has either of you looked at that paper and thought, "I really think it should have that in it"?

Michael Birtwistle: Privacy is not in it, but I think that the team behind the White Paper took the view that that is already fairly comprehensively covered by data protection law. There are questions about not the high-level principles, but whether all the instructions to regulators that the principles represent will necessarily be deliverable by those regulators under their existing powers and remits.

Dr Oswald: The White Paper talks a lot about focusing on use, not on the technology. That is fair enough up to a certain point, but I think that there are certain technologies or implementations where we have to look at the tech itself such as emotion AI, which comes from a very dubious scientific background—similar to the polygraph—

Katherine Fletcher: We all love a bit of phrenology.

Dr Oswald: Exactly. The scientific validity is very dubious there, and the ICO said that. At some point, we have to say, "Do we just not want to use that?", especially in criminal justice where serious consequences can occur.

Q530 **Katherine Fletcher:** Right, so the ability to exclude low evidence or no evidence-based technologies could be included in a future White Paper.

Dr Oswald: Yes.

Chair: Thank you, Katherine. Next, Stephen Metcalfe, then finally Graham Stringer.

Q531 **Stephen Metcalfe:** Michael, I have a quick question about a statement made by the Ada Lovelace Institute on the use or importance of standards to create an AI governance framework. Will you explain exactly what that means, and where the benefits would be?

Michael Birtwistle: Standards will have a very important role in any potential regulatory regime for AI—

Q532 **Stephen Metcalfe:** Before we go on, will you define what you mean by standards? There is always a danger that everyone is different.

Michael Birtwistle: I think we are referring to process standards, technical standards, the sorts of things that standard-setting bodies such as the BSI in the UK, the ISO or IEEE—those sorts of bodies—set. It is about using those as a way of helping to govern and standardise the way in which AI is developed as a process. You have horizontal process standards and you can also have vertical standards, which are, for example, used very commonly in medical devices regulation.

Q533 **Stephen Metcalfe:** Thank you. How do you think this will work within AI? What sort of standards would you like to introduce, or do you think would work?

Michael Birtwistle: I think the White Paper sets out a very helpful vision of the role of standards and assurance techniques, which are things such as audit, that will help to support the way that different actors in the ecosystem working with AI know what is right to do. Also, from a regulatory perspective, regulators are able to use those, potentially, to help to understand whether companies look like they are doing the right sort of things to make sure that you get the right outcomes from AI development.

I think there are three interesting questions around the use of standards for AI. The first is, who gets to participate in setting them? We are doing a lot of work in Europe at the moment around the fact that these are typically set largely by industry, and there is perhaps a degree of industry capture in who gets to set them.

Secondly, there is a question about regulator maturity. Some regulators are really well equipped to use them and use them as standard—some very much less so. But they could be a useful tool for regulators.

Lastly, there is the question of the international picture. The US and Brussels are obviously taking a very active interest in the standards that will support the development of AI, and the UK to some extent will be competing for influence over what those standards become.

Q534 **Stephen Metcalfe:** Have you made any suggestion about what those standards should be—any practical steps that we could log?

Michael Birtwistle: I do not think we have specified standards. I think we have specified, within biometrics governance for example, a role for the function to set those standards or identify those standards. Those standards are typically developed by those standards-setting organisations. We would want to see the principles in the White Paper supported by those standards, but that I believe is work in train in Government.

Q535 **Stephen Metcalfe:** Are you aware of anyone who has been so bold as to make any suggestions about what standards AI should be governed against?

Michael Birtwistle: Yes, there is lots of active work, and the UK has set up a standards hub that is working with international partners on which are most appropriate. I believe there is work ongoing in Government to understand specifically which sets of standards can support the individual principles.

Q536 **Graham Stringer:** The UK looks as though it is diverging from the EU on a regulatory governance framework for AI. Is it better to govern AI at a global level, or is it better if each country goes its own way?

Michael Birtwistle: The Ada Lovelace Institute supports the interoperability of these AI technologies. That is going to be very important for ensuring that their benefits are distributed equitably and that they support trade between nations. I think there is a question about which bits of it you govern at what level. There is potentially a role at the international level to look particularly at the urgency surrounding general-purpose AI systems—foundation models, as they are sometimes called—to look at some of the emerging capabilities that those systems have and to understand whether an international approach will be most effective at mitigating those quite wide-ranging risks from very powerful systems. But I think it is also right to have a context-specific approach for our country for more detailed applications, as Marion and I have already said.

Dr Oswald: If I could add to that. It is important that, in the criminal justice context, we take account of those applications at the EU level that have been identified as the highest risk. That will have implications for trade, criminal justice co-operation and so on. We need to make a decision on how we deal with those ourselves.

Q537 **Graham Stringer:** If you could have one recommendation that the Government would agree to implement on their future policy regulations of AI, what would it be?

Dr Oswald: From a policing perspective, we need a statutory framework that includes an independent advisory function that looks at AI deployment from end to end—from start to finish—and carries that rolling review.

Graham Stringer: Okay. Michael?

Michael Birtwistle: From a policing perspective, I have already hammered on about biometrics governance. If we are looking at the White

Paper, the one thing that we would like to see clearer commitment around is the scale of resourcing that AI governance will require. It is useful to look at AI through the lenses of infrastructure and of safety. If we look at the level of resourcing that we put into regulators that look at those domains, we will see that the Civil Aviation Authority has £140 million revenue a year, the Office of Nuclear Regulation £90 million a year, and the Office of Rail and Road £36 million a year. That is probably the order of magnitude of the resourcing that you should be looking at for governing a general purpose technology of the significance of AI, regardless of whether you deliver that on a centralised or distributive basis through existing regulators.

Chair: Dawn had a supplementary question.

Q538 **Dawn Butler:** Thank you both for your evidence today. I just want to follow up on the question that Katherine was asking in relation to privacy.

Michael, you said that privacy is not currently in the White Paper, but it is covered by the Data Protection Act 2018. I think it is also covered by the Human Rights Act 1998 and the European convention on human rights. As we know, the Government want to scrap the Human Rights Act. Do you then think that, as a safeguard, privacy should be in the White Paper?

Michael Birtwistle: That is a difficult question to answer hypothetically. The research that we have published around, for example, facial recognition technology demonstrates that there is a great deal of public appetite for privacy to be balanced correctly against the security benefits of some of these technologies. Regardless of the source of that protection, I think it would be important for privacy to be an integral part of the balance of interests that you consider when you are deploying artificial intelligence.

Dawn Butler: Thank you.

Dr Oswald: If I could add, proportionality is that key concept within human rights legislation and it is embedded everywhere in the review of technology in policing and national security, including around investigatory powers. However, it is a really difficult assessment to make as to whether or not something is proportional. It is really easy to say, "Well, we need to do this for policing purposes and therefore it is proportionate," but that means that you are not assessing whether something employs the least intrusive means.

I am the joint author of a recent report that was published yesterday by the Alan Turing Institute into how we could better use different factors to do that proportionality assessment when we are talking about AI and automated analysis. That is the sort of thing that could be followed up by the White Paper to give people better guidance as to how to do that proportionality assessment. I think it is crucial to keep that within our framework because it is fundamental in criminal justice.

Dawn Butler: Thank you both.

Chair: I thank Dr Oswald and Mr Birtwistle for their evidence today and say to all our witnesses that it has been an insightful session into the very important application of AI. That concludes the meeting of the Committee.