Telecommunications and Media Forum: Artificial Intelligence and the future Digital Single Market

Keynote speech on privacy, data protection and cyber security in the era of AI

24 April 2018

Giovanni Buttarelli

Many thanks for your invitation this event.

The questions raised by AI are myriad, many legal but more importantly ethical. Indeed as the big subject of the World Economic Forum in Davos in January, privacy was flagged as the biggest concern surrounding the development of AI systems.

I would like to explain why this might be, and what the EU should do about it. Good regulation, in our view, is only part of the answer. And we still have a lot to do on that score, even with the GDPR becoming fully applicable in one month's time.

We know that 2018 will go down as the year of data protection not only because of the GDPR, but also for less congenial reasons.

This year practices have been coming to light which contradict the most basic principles of not only data protection but basic respect for people who have been goaded and cajoled into putting more and more of their lives online. We are seeing just the tip of the iceberg.

So I will make three propositions.

First, that self-regulation in the era of AI must play a role but it urgently requires the establishment of some ethical parameters.

Second, regulation - like the GDPR and rules on confidentiality of communications - is indispensable and needs to be completed.

Third, regulation of how data is processed and how communications services are provided is not enough. We have to address the question of structural imbalances and unfairness in digital markets.

First, on the role of self-regulation. Laws cannot fix all problems. You might be forgiven for thinking that the GDPR was attempting to fix everything, given how long and prescriptive it is.

The last few years have demonstrated that digital markets cannot be left entirely to their own devices. Doubts surfaced with a string of high profile data breaches, like the Ashley Maddison incident in 2015, controversial not so much in the volume of data but the sensitivity of the type of data. It is culminating now with the Facebook / Cambridge Analytica case.

This has shone a public spotlight on the wider problems with the digital information ecosystem.

It is interesting how the arc of this news story has moved beyond being a local scandal about 'fake news' to a symptom of approaches to people's personal information on a mass scale. Approaches which now seem unlawful, and almost certainly unethical.

Companies therefore should be able to take calculated risks with new products, in the light of honest assessments of the likely impact on people, and an analysis of possible unexpected consequences.

That is why the EU's approach emphasises the principle of accountability.

In other words, if you derive value from processing personal data, then you must give account to the people affected, as well as to the regulators whose job is to oversee compliance.

That is why the obligations of the controller are at the centre of the GDPR, along with provisions like codes of conduct and certification for demonstrating compliance.

Accountability is the biggest challenge when we think of AI. AI is now the most fashionable pretext for collecting data.

It requires personal data on a huge scale, at least until - like Alpha Go - it becomes intelligent enough to teach itself. Bias in AI systems is therefore inherent to the training data rather than the algorithms.

The volume and sensitivity of data processed for developing AI, and the decisions which come out of AI, raise basic questions of accountability: if harm is caused by an AI system which had been developed and delivered value for a profit-seeking company - who should be held responsible for that action?

It is one thing to apply AI to the harmless abstractions of a board game, or to chatbots in a call centre. It is another to apply it to supporting the mentally ill or educating children. Still another level of potential harm for autonomous vehicles.

There is also a growing trend of trying to vindicate, post hoc, rampant data collection by identifying potential 'philanthropic' uses for the data. In fact, we need to clarify the public interest – if necessary through democratically adopted legislation – for using big data and AI.

We already have basic ethical framework in place.

Industry is very familiar with principles like Know Your Customer and Due Diligence which are not legal requirements but rather cultural norms for accountable organisations.

What I find most remarkable is that during the timeline of events leading to the current scandal, ethical questions have seemed to be entirely absent from decision-making.

For instance, the word 'ethics' appeared only once in 10 hours of evidence before the US Congress committee hearings this month. The perverse incentive in digital markets to treat people like sources of data has to be remedied. Ethics will be playing an increasing role.

Which brings me to my second point, on completing the necessary regulatory framework.

Older sectors like telecoms, broadcast and print media appreciate the historical imbalances in regulation of sectors. But for two decades tech has been largely outside the reach of regulation. They have been allowed to move fast and break things.

Now there are a lot of broken things which need to be mended, and we need to guard against future breakages.

In reality, regulation rarely chokes innovation *per se*. Rather, it provides a new set of market incentives.

In the gaps between obligations and prohibited practices, there is a vast hinterland of possibility. Good regulation steers innovation away from potentially harmful innovation and into areas of this hinterland where society can benefit. But this is also the space where bad habits have sprung up. And so now the dominant business model for web-based services requires maximum data collection, tracking of behaviour, fostering addiction, experimenting with stimuli to elicit reactions.

The result has been an assumption that, in order to be profitable, you must collect as much personal data as possible, by whatever means possible, and then try to find ways to monetise that data.

Only now are people at last realising that they are being tracked across the internet, whether or not they are logged in or have an account with one of the major tech platforms. You cannot escape this.

To paraphrase the Eagles' Hotel California, you can log out of the internet anytime you like, but you can never leave. Commercial tracking and targeting is acceptable in some cases. But in the last year or so, the scandal has spread to the civic space.

Our political system, even more than the economy, is based on notion of free choice at elections. Where even this seems to be in jeopardy, there is a clear need for more rigorous and coherent enforcement.

The GDPR is essentially a continuation of the EU's 1995 rules. Will we notice big changes on the ground? Promotional activity surrounding GDPR compliance now abounds. How much of this is reflected in genuine safeguards for individuals?

Already in the last week we have seen one major company deciding to move around its operations in order to avoid applying GDPR to data about people outside the EU. We shall see.

On 25 May, data protection regulators will reconvene as a new legal entity, the EDPB, and start working to bring about a more consistent and effective enforcement of rules.

Meanwhile, there is still a lot of confusion about the necessity and rationale of revised ePrivacy rules.

Do we need for rules at all, now that we have the GDPR?

There is a big technical reason why we do - it is not fair to expect an enormous and valuable sector of the EU economy to be subject to harmonised data protection rules, on the one hand, and national ePrivacy rules under Directive 2002/58, on the other.

Certainly we could all do better in explaining how GDPR and ePrivacy mesh together. The GDPR is a big piece of the jigsaw and a massive achievement. But it is basically a set of minimum standards for handling personal information. It does not specifically address the sector of the economy which provide communications services.

Communications are meant to be, by default, confidential. That is what people expect and that is what the EU fundamental right to privacy requires in Article 7 of the Charter. ePrivacy must now to cover services beyond traditional telecommunications services and network providers.

If properly reformed ePrivacy will go a long way to stopping constant snooping on people's communications via services and apps. It will stop companies forcing people to accept being monitored ('tracking walls') in exchange for accessing content online.

So again, the EU is trying to change market incentives, encourage innovation so that access to information on the internet does not depend on being watched all the time. And again, like the GDPR, it is not a panacea. It is one part of the structural remedies which we are trying to bring to a broken system.

There are many strange myths surrounding the proposed ePrivacy regulation. One of the most curious of these myths says that requiring consent for practices which most intrude into people's privacy will drive more business to the dominant players. In fact studies show that requiring meaningful consent, with tough enforcement, will hit unsustainable practices hard.

DPAs are getting ready to do this.

Consent like any legal obligation is susceptible to becoming a tick-box exercise, where the spirit of the law is flouted even as the letter of the law is apparently respected. But consent is central to data protection and all 'free' transactions in society because it is a means for giving people control over processes decisions which affect them.

Empowerment is not enough on its own. Accountability in the GDPR is perhaps the more important principle. But consent and individual control are indispensable when it comes to interference with the most intimate space - sensitive data and communications.

That is why legitimate interest as a legal basis for data processing is not appropriate in the case of communication metadata and location data. It would create legal uncertainty, especially for data processors.

People are demanding stronger privacy protections not potential loopholes.

Of course, a number of sticking points remain with the ePrivacy proposal.

I know for instance that the provisions on machine-to-machine communication remain problematic for many, especially in the telecoms sector. Distinguishing machine-to-machine communications is misleading and a recipe for legal uncertainty. Communications which could contain sensitive information should be confidential - it doesn't matter whether the sender and or recipient is a human or no.

To carve out such a technical exemption would go against the objective of increasing trust in communications services. There may be some scope for sensible solutions which preserve the key value of the instrument.

We support a balanced solution between innovative reuse of personal data, subject to user's control with a high level of protection of fundamental rights.

I have heard it said that the ePrivacy regulation would not have prevented the Facebook Cambridge Analytica case. Perhaps not. Data protection is part of the solution.

So my third point is that the EU must use all tools, in particular antitrust, to decentralise the internet to give people more freedom and choice online.

The current controversy highlights systemic issues - it is not a problem with one or two rogue companies here and there. The digital dividend is not being fairly shared.

Markets have become so concentrated that a handful of players are able to determine terms and conditions which competitors are expected to emulate and which individuals cannot negotiate or contest.

They control the flow of information across the digital ecosystem. Advertisers and publishers depend on them but cannot evaluate whether they are getting a good deal or not. The picture is replicated in the market for AI.

The annual value of mergers and acquisitions in artificial intelligence companies has increased by over 500% since 2013.

It may be that some companies become so powerful that they become an existential threat to democracy – that is the root of antitrust, both in the United States and Europe, and especially post war Germany.

It does not matter whether or not the companies in question have benign, noble intentions. In such cases, the potential for harm becomes more important than the *motivation* to harm.

Look for instance at the persecution of the Rohingyas in Myanmar and the role social media has played in this. This was the message of our opinion last month on personal data and online manipulation.

We cannot be in a position where a company is in effect too big to comply, for instance, with a request from a data subject for information held on him or her.

Our initiative of the Digital Clearinghouse has given different regulators a confidential forum in which to consider common concerns.

There is a recognition of that a coordinated approach is actually not an option.

Therefore in the light of recent scandals, we will open up this forum to regulators responsible for audio visual services and elections and host an event early next year before the European Parliament Elections.

So to conclude.

First, there is a clear role for self-regulation, but we need a conversation about ethical underpinnings.

The problem with the explosive digitisation of society has been, at worse its *a*-morality, or at best naiveté, about what constitutes the social good. Our conference at the end of this year will help, I hope, plot a course towards a new digital ethics.

Second, we need to complete the legal framework by urgently adopting privacy rules for electronic communications.

Finally, we also need coherent enforcement across domains, the only way to begin to address the systemic problems we are facing.

Thanks for listening.