ECIS European Committee
for Interoperable Systems

ECIS ivzw-aisbl
Louizalaan 65 avenue Louise Box 2
B-1050 Brussels, Belgium

T/F +32 (0)2 706 24 15
info@ecis.eu
www.ecis.eu

ECIS event: "The Future of AI in Europe"
29 November 2017
Summary of presentations

Hannah Crowther

The GDPR was in essence created to address challenges posed by new technologies. However, some aspects of the law continue to create challenges for those technologies, such as for instance those concerning automated decision-making (and data transfers). AI is a very broad concept, and so it can be helpful to analyse it on a sliding scale:

- Using AI to learn something, make discoveries: the GDPR is permissive here as long as you have transparency.
- Using AI to predict or profile people: GDPR makes life slightly more difficult for this. There is a need to think about balancing rights.
- Using AI for automated decision-making: Whilst we interpret this as a prohibition, not a right to object, the GDPR makes life very difficult in this situation.

What is a decision? When does it have a significant effect? Are we talking about all automated processing or only when machines exercise discretion – and what is the difference? It is just code? When you have code responding to an input, there is a straightforward decision. Thus, it does not really feel like the kind of decision we are talking about when we are talking about AI – it is more of a fact, a programmed input. The prohibition on automated decision-making should not be a prohibition on automated processing.

When we talk about human intervention, does the human being have to second guess the decision? Is there discretion around the decision? One concern with maintaining a *prohibition* (rather than a right to object) on automated decision-making is that it might lead to a lowered standard of protection, with people trying to shoe-horn out of the "automated decision" rules altogether, instead of focusing on including safeguards such as human intervention.

Dimitra Kamarinou

When looking at lawful processing under the GDPR, you have to look at the whole of the decision-making process (collection of data, model development, decision-making stages). There are derogations for wholly automated decisions: when authorised by law etc. However, even if allowed, data subjects have a right to obtain human intervention on the part of the controller, and they have the right to express their point of view and to contest the decision.

In terms of fair processing under the GDPR, data controllers have obligations to use appropriate mathematical or statistical procedures for profiling, implement technical and organisational measures (to avoid discriminatory effects). However, it seems difficult to achieve this: some algorithmic models work better through access to large amounts of data – and this data can be biased both directly and indirectly.

The principle of data minimisation according to which data controllers only have to process only the data that is strictly necessary for a specific purpose previously determined by the data controller, might cause problems as data is inherent in the nature of machine learning.

In terms of transparent processing, some of the AI may work in a non-transparent way because the underlying algorithms are designed as such.

In terms of capability of AI systems to explain a decision, what do we mean by "meaningful information about the logic involved" in a decision? It seems clear this should be assessed from a data subject's perspective.

There is still a debate around whether an explicit or implicit right to explanation exists under the Regulation. A European Parliament resolution said that as an ethical principle, 'it should always be possible to supply the rationale behind any decision taken with the aid of AI that can have a substantive impact on people's lives' – 'advanced robots should be equipped with black boxes to record data on every transaction carried out by the machine including the logic that contributed to its decisions' (paragraph 12 of the EU Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics) (this seems to be more explicit than in the Regulation).

From the compliance side, data controllers have to carry out risk assessments (Data protection impact assessments) for new technology, when they process sensitive data on a large scale, when they monitor public spaces on a large scale (*e.g*., with facial recognition). The risk assessment must describe processes, potential risks etc. The UK Information Commissioner's Office (ICO) has suggested that Companies may also have to set up ethics boards.

The potential for fines if data controllers or processors breach fundamental rights might cause reluctance to use some data or some technologies.

The European Parliament accepts that in the long term, there is a possibility that AI may surpass human intelligence – an interesting possibility is the prospect of AI overcoming certain key limitations of human decision-making and providing us with decisions that are demonstrably fair.

Paul Nemitz

The word "ethics" is over used. Programmers have to come to the table and tell the Commission what is happening in order for the Commission to understand what needs to be done. But what these programmers/companies tell the Commission has to be true.

Why do we have this big debate about AI? Is it a mere buzz word or is there really a challenge that we really have to look into? A bit of both it seems. If the potential long term impact of AI could be problematic, we have to invest into impact assessments to avoid long term damage.

I think we need to clarify the terminology: automation is nothing new. But this is different from autonomous decision-making. The latter creates a causality which cannot be linked to the creator of the program. In the "right to be forgotten" case, Google argued it was not a responsible controller because the algorithm was giving the answer. This was an attempt to get out of responsibility – but it was held to be a responsible processor. However, in other areas of the law such as the law regarding product liability, the answer may not be as straightforward.

In relation to next year's Communication, here is what I believe we need to do. Where we have a technology that becomes autonomous and becomes like an infrastructure (an input for everything) – including everything concerning rights of individuals – we need to incorporate some basic principles in the development of the technology. This is what I call responsibility by design – design it to comply with democracy, fundamental rights and the rule of law. From this, it appears clear that individuals need to know whether they are speaking to a human being or a machine – and there should be a binding rule for this.

Concerning the relationship between law and ethics, one school says we need law, the other one says we do not want binding law but want ethical rules. We need to know which challenges needs ethics and which challenges need law. I suggest we have a point by point debate. We also need a public catalogue to know where the state is using AI. For instance, we can imagine a scenario where a private company is deciding on parole and arguing that how it does so constitutes a business secret.

In relation to codes of conduct, the relevant questions are: (i) how do they relate to the principle of democracy? and (ii) are they legitimate if not made by people who are elected?

Turning now to the existing law, the important word in article 15 of the GDPR is the word "meaningful". The Article 29 group paper on automated decision-making, which gives examples of how data controllers should respond to requests, is a bit thin. My advice is to be more forthcoming – the word meaningful has to be understood generally. You have to interpret this in light of the charter of fundamental rights and notably the principle of self-determination, with the individual being at the centre. You also have to look at internal market law. I believe that judges will look at the charter of fundamental rights and interpret this in a way that is friendly to individuals.

Human beings are subjects, not objects – we cannot accept that we are controlled by a machine and this is more important than the economic aspects.

Thomas Vinje

As an antitrust lawyer, I will touch upon some of the antitrust issues that AI may bring about. But let me begin by saying that AI is an opportunity we should embrace – we should allow Europe to innovate and we therefore need to be careful not to over regulate.

AI may raise some competition law concerns. Indeed, AI may enable collusive conduct. It may be feasible for competitors in a market to program their AI algorithms in a way that causes them to engage in unlawful collusive conduct, for instance by determining that the most profitable way to respond to a price increase by a competitor would be to increase one's own price. However, collusive behaviour is not the only thing we need to worry about. One can imagine an abuse of dominance case where an AI machine figures out it can for instance maximise revenue by using a strategy of loyalty discounts, which can be unlawful in certain circumstances.

Even if there are some competition risks associated with AI, AI is generally a good thing and is beneficial to consumers. To ensure that consumers benefit from AI, it is crucial that companies can freely compete and innovate in AI technologies. I suggest we take a look at how competition law could play a role in ensuring that companies are not anti-competitively prevented from competing and innovating in AI.

So what feeds AI? Data. Access to data has a crucial role in ensuring companies can freely innovate in AI. Indeed, the best algorithms out there cannot achieve good results if they have insufficient data from which to learn and improve on their results. A particularly salient illustration of this is provided by online search services.

Because of the particular machine-learning nature of a search engine, it is query scale – not technology – that is the primary driver of search engine profitability and competitiveness. Search algorithms learn from user queries and how users interact with search results, and the higher the number of queries (in particular the so-called "long tail" queries), the more relevant results the search engine will be able to show to users.

This helps explain why barriers to entry can be so high in a search engine market where one company has an overwhelming advantage in scale of queries. If dominance is achieved via scale in data, a company has an incentive to undertake conduct to preserve its position and prevent others from obtaining requisite scale to compete – and such conduct can be anticompetitive.

Such behaviour is of great concern for AI as unlawful foreclosure of competitors will inevitably lead to less competition and hence less innovation in AI. Competition authorities should be carefully looking at this.