Policy Brief

A European approach to regulating AI
The Commission's AI strategy

5 November 2020

Maarja Kask, Policy Fellow

The European Commission's AI strategy was released just a few short weeks before the global pandemic hit Europe. In her policy brief, Maarja Kask argues that AI can only be expected to help us through this pandemic and on a path of economic recovery if the upcoming legislation resolves the issues evident in the white paper. Not only does the Commission need to explain what it means when it talks about AI, but also be clear about its investment plan. The member states will also need to pull their weight when it comes to funding for the strategy to succeed.

Almost immediately after the European Commission set out its new digital strategy on 19 February, the COVID-19 pandemic completely disrupted our normal way of life and, among many other things, altered the digital landscape. It raised questions about European digital infrastructure and co-operation like the sharing of medical data. It also shifted the focus of the artificial intelligence (AI) debate to AI in healthcare. Digital solutions like AI are expected to help tackle the pandemic. Moreover, the reliability of the digital infrastructure is now much more important and prevalent with people not only working from home but spending more time on the internet and using internet-connected devices. All of this once again underscores how crucial it is to regulate Europe's digital sphere in a manner that protects Europeans while not hampering innovation, one that supports the uptake of digital technologies without enlarging inequalities.

This policy brief outlines the main elements of the Commission's new digital strategy and assesses its key components. It argues that the general approach to AI regulation is well taken but lacks a focused definition of what AI is. The Achilles' heel of the strategy will likely prove to be insufficient funding. As the EU will be unable to make much of a difference due to the small amounts allocated to the policy area in the next long-term budget (MFF), member states will need to raise their own financial contributions. The Commission, meanwhile, should increase transparency about its own investment targets and how it aims to fulfil them.

New policies and legislation

"A Europe fit for the digital age" is one of the six political priorities of Ursula von der Leyen's Commission. To give meaning to this slogan, it released three documents in February 2020, which together constitute its new digital strategy. The three papers are highly interlinked, with the communication "A European strategy for data" supporting the white paper "On Artificial Intelligence – A European approach to excellence and trust". The communication "Shaping Europe's digital future" binds the two together and gives an overview of the upcoming actions and plans that will be put in motion to provide a framework for the digital sector. Further, a report on the safety and liability implications of Artificial Intelligence, Internet of Things and robotics was released, focusing on the legal liabilities concerning digital solutions.

Legislation on AI was initially scheduled to emerge in the first 100 days of the Commission; however, the regulation was 'downgraded' to a white paper, accompanied by a public consultation process amid speculation about a regulation coming out by the end of 2020. The task might have proved to be more complicated than expected, as the Commission's updated working plan now states that the legislation on AI will come out next year.

In the AI strategy, the Commission outlines its plan for a proper regulatory framework for AI. The Commission has opted for a two-pronged approach: a range of initiatives to support the uptake of AI in Europe is called the 'ecosystem of excellence', while the regulatory framework is an 'ecosystem of trust'.

The policy framework – the ecosystem of excellence – contains different initiatives to increase and support the uptake of AI technologies in Europe, as well as to align efforts at European and national level. The assumption behind the Commission's strategy is that there will be a paradigm shift from consumer data held in cloud services to that of industry, business, and public sector data stored locally on devices. As Europe has a strong industry and offers corresponding industrial services, the hope is that industrial data will be Europe's stronghold in the global data economy. The eight fields of action foreseen by the Commission are:

- Working with the member states: a number of joint actions for cooperation, including an update on the Coordinated Plan on AI

- Focusing the efforts of the research and innovation community: creating centres of excellence in research and innovation for AI that attract investment and talent (may include a new legal instrument)

- Skills: a reinforced Skills Agenda to overcome shortages

- Focus on SMEs: help and access to financing for SMEs keen to deploy AI technologies

- Partnership with the private sector: setting up a PPS in AI, data, and robotics

- Promoting the adoption of AI by the public sector: establishing "sector dialogues" to facilitate the adoption of AI in areas such as healthcare, administration, and transportation

- Securing access to data and computing infrastructures: investment in computing technologies and data infrastructures

- International aspects: continuing work on international cooperation on guidelines and standards for AI

The regulatory framework – the ecosystem of trust – contains legal guidelines for AI technology. The Commission's overall expectation for 'AI made in Europe' is one of a trustworthy, human-centric AI. It emphasises that for new technologies to be taken up, citizens need to be able to trust them. As per the recommendations of the High-Level Expert Group on AI, this includes seven key requirements that all AI applications will have to meet:

- Human agency and oversight: user wellbeing at the centre, with humans in control through different in- and on-the-loop measures

- Technical robustness and safety: reliable and resilient even in case of errors, security-by-design mechanisms

- Privacy and data governance: data protection to be guaranteed in all life cycles of the system, all actions with data to be documented

- Transparency: decisions made by the AI system should be traceable and explainable, it should be made clear whenever AI systems are being used

- Diversity, non-discrimination and fairness: the biases in datasets should be considered when developing AI systems, design and development teams should be diverse and inclusive

- Societal and environmental well-being: AI systems should incorporate ecological responsibility, AI should have a positive societal impact

- Accountability: AI systems should be externally and internally auditable, negative impacts of any systems should be documented and minimised

Beyond these general requirements for all applications, the Commission is advocating for a risk-based approach. This means that some AI applications should be classified as high-risk and regulated accordingly while others should be regulated less strictly as they are considered low risk. According to the white paper, an AI application should be considered high-risk when both the sector and usage involve significant risks to safety, consumer rights or fundamental rights. Sectors concerned here are, for example, healthcare and transport. The paper states that these sectors should be exhaustively listed in the regulation. Usages rendering an application high-risk depend on its impact upon affected parties. This includes, for example, applications with legal consequences.

In exceptional instances, an application itself could be considered high-risk. Examples include systems that endanger human life, those used in national security and law enforcement, systems used in providing services by a public body, or those used in providing services of public interest. One of these exceptions is also remote biometric identification, which is considered one of the most dangerous forms of AI application, posing a significant threat to privacy. AI systems considered high-risk will have to fulfil additional requirements concerning:

- Training data

- Data and record-keeping

- Information to be provided

- Robustness and accuracy

- Human oversight

- Specific requirements for certain particular AI applications, such as those used for purposes of remote biometric identification

All of these requirements are to be described in detail in the regulation, to ensure that actors bound by them can comply with them. The Commission suggests a "prior conformity assessment" to verify compliance.

For applications that are not high-risk, the Commission suggests a voluntary labelling scheme. As these applications do not have to meet any additional requirements, operators may want to add a label to signal the product's safety. To earn this label, the application should fulfil requirements specific to low-risk applications. These requirements would be developed under the scheme. While joining the scheme would be voluntary, the operator of the product would be bound by

the rules after gaining the label and thus be subject to enforcement.

The scope of regulation needs to be defined

Regulating AI is not an easy task – developing AI technologies is crucial for the digital economy, yet these technologies may generate significant discrimination and inequality issues (or reproduce the ones created by humans) if not regulated properly. At the same time, these technologies evolve rapidly so any regulation must be able to anticipate changes and leave room for possible future developments. Moreover, AI is a relatively vague concept; there is no general consensus on which digital technologies count as AI and which do not.

The lack of a proper definition is a significant issue for creating meaningful regulation. In the white paper, AI is defined as being "a collection of technologies that combine data, algorithms and computing power". This definition, however, could be applied to almost any software, thus creating legal uncertainty in determining which applications the legislation affects. Two further definitions are given in subsequent footnotes, one of them the definition given by the HLEG on AI: "Artificial Intelligence (AI) systems are software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions" [sic].

To make matters even more complicated, definitions of AI differ according to their purpose: technical definitions do not fill the same purpose as legal definitions and are thus unsuitable as a basis for regulation – the definition of AI given by the HLEG is an example of this. The Commission will have to make sure the definition follows the principles of a proper legal definition:

- Inclusive: the goals of regulation must not over- or under-include

- Precise: it should be clear which case falls under the definition and which does not

- Comprehensive: the definition should be understandable by those who are regulated

- Practicable: legal professionals should be able to easily determine whether a case falls under the definition

- Permanent: the need for continued legal updating should be avoided

Given the nature of the technology, the principle of permanence may be difficult to follow. The legislation may have to be amended periodically. Precision and inclusiveness, however, are crucial. The Commission will have to take care not to define AI simply through different techniques (e.g. machine learning, computer vision), but to take into account the technology's effects and outcomes. Moreover, one has to consider different stages in the lifetime of the AI system, and ensure that different services conducted by AI systems and the outcomes from these services fall under the scope of the definition. This holds true even after the act of service itself has long passed (we can imagine, for example, a situation where the consequences of a decision made by AI only appear decades later). Overall, the Commission should not insist on a simple, one-sentence definition, as this would not amount to a sufficient basis for regulation, but must include a thorough explanation of the regulation's scope, detailing what is considered to be AI under it.

Applications of AI in different sectors are extremely varied and deploying AI in these different sectors results in very different outcomes and accompanying problems, leading to the conclusion that a one-size-fits-all approach is not suitable for AI regulation. For example, the protection of data subjects is extremely important in healthcare, while in industry or mobility it may be less so. At the same time, standardisation has also proved itself to be a difficult task. Because of this, over-arching regulation may result in curbing innovation, while failing to provide protection where necessary. Taking all of that into account, it is best to create very sector-specific regulation.

The European Commission has opted for a risk-based and sector-specific regulation. With this kind of regulation, all regulatory actions will be targeted to specific sectors and to the use of the system. Nonetheless, the Commission's expectations for 'AI made in Europe' are high and while sectoral-based legislation strikes the right balance, some of the requirements foreseen for all applications may not be necessary to apply in every circumstance. This concerns in particular the transparency and human oversight requirements.

Some AI applications can be extremely simple: for example, the junk or spam filter of an email programme – requiring human oversight for these applications would defeat their purpose. While transparency is something to strive for, its absence does not necessarily cause risks. When it comes to decision-making, transparency is clearly important in decisions like handing out loans, for example. But when AI is significantly better at detecting cancer than human doctors, we do not necessarily need to know why the AI system decides to categorise one x-ray as cancer and not the other. This might, however, become important in the case of liability issues, for example, if a medical procedure based on a decision by AI results in injury. Thus, transparency might become important after the AI has already "done its job"; in other words, different levels of transparency might be necessary throughout the lifecycle of the AI. The Commission thus needs to specify which requirements are truly necessary for all AI applications, and in which stages of application.

As the Commission moves further with the regulation, it must define its scope. This means, firstly, giving a clear legally applicable definition of AI that meets the appropriate requirements. Secondly, it is important that the requirements given by the Commission for all AI products are specified, and the question "Which requirements need to be fulfilled in which stages of the system's lifecycle?" can be answered clearly.

Funding matters

In the AI white paper, the Commission states that "Europe can combine its technological and industrial strengths with a high-quality digital infrastructure and a regulatory framework [...] to become a global leader in innovation in the data economy and its application". At the same time, new technologies are expected to help fight climate change and other pressing issues. However, this digital infrastructure needs to be developed further in order to stay up-to-date and be able to effectively support solve current issues.

The 'ecosystem of excellence' – Europe's deployment of new technologies like AI – is also highly interconnected with the capacity to conduct research and to innovate. For example, the speed and outcome of AI systems depend on how much processing power the computer has. To make considerable progress in increasing processing power in computers, advances in quantum computing are necessary. For these to come to fruition, significant efforts in research and innovation are needed. And Europe needs to fund these efforts.

The European Commission's goal is to reach more than 20 billion euro of total investment per year in AI alone. However, in 2016, 3.2 billion euro were invested in AI in Europe, compared to €12.1bn in North America. In the MFF, we can expect investment in AI to be found in the Horizon Europe (HE) research programme and the Digital Europe Programme (DEP). From the latest MFF (2021-2027), the DEP is allocated just €8.2bn – covering the whole seven years. Of the DEP, 2.2bn euro will be allocated to artificial intelligence and €2.4bn for supercomputing, which will support the uptake of AI. It is difficult to estimate the amount of investment in AI in the HE programme but, looking at Horizon 2020, we see that €1.5bn was invested in the period 2018-2020. Taking that into account – and adding a very optimistic percentage – we can estimate roughly 5bn of the total 80bn euro in HE would be allocated to AI over the seven years. This means that the DEP and HE programmes together account for approximately 10bn euro, giving €1.4bn per year. Taking as base value the Commission's recommendation of 20bn euro per year, this leaves us far off target, with member states having to cover at least €18.6bn annually from public and private resources.

When it comes to the Commission's own investment in AI, the numbers are not very transparent. As previously mentioned, there are no clear estimates for AI investment in Horizon 2020, nor are they very clear for Horizon Europe. For HE, an investment of 700 million euro is mentioned in the policy and investment recommendations report by HLEG on AI. This, however, seems

exceedingly low and gives no reference as to where this number originates from. On the Commission's homepage for artificial intelligence, we find the number of 1.5 billion euro invested in Horizon 2020 between 2018 and 2020 and the target of 20 billion euro of total investment per year. Unfortunately, this leaves us in the dark about how much the Commission itself expects to invest in AI in the upcoming years. What is the Commission's real AI investment target for Horizon Europe? How much of the €20bn annual investment target does the Commission expect to take upon itself? How do the AI investments on HE and DEP get allocated? These questions appear unanswered. The Commission needs to become significantly more transparent about its funding targets as well as how it aims to fulfil these targets.

Member state investments should also increase. However, as seen in the debates about the MFF, member states are not fighting for more investment in future-oriented programmes like HE or DEP, but rather for base programmes like the Common Agricultural Policy and cohesion policy. Nonetheless, if unwilling to contribute through the MFF, they will have to draw on their national budgets. The USA and China are stepping up their efforts in AI investment and, to stay competitive, Europe needs to raise its own game. European countries may want to consider bundling investments for greater effect. For example, in 2019, Germany and France agreed on creating a common research and innovation network for AI.

Conclusion

The European Commission is on the right path by creating regulation. Europe needs a common regulatory framework and pooling of investments to avoid single market fragmentation and to develop the best possible data infrastructure and AI 'made in Europe'. For the policy to be effective, the regulation's scope needs to be clear and the definition of AI should meet the requirements of a legal definition. If Europe is to seize all the possibilities that the digital sector offers, efficient investment is crucial. It needs investment in infrastructure but, beyond that, Europe's digital ambitions cannot succeed without investment in research, innovation and deployment all over the continent. For that, the member states need to do their part, while the Commission must be entirely transparent about its own planned investments.

This publication is part of the research project "Repair and Prepare", a joint project of the Bertelsmann Stiftung and the Jacques Delors Centre.
For more information, please visit www.str: www.strengtheningeurope.eu.

Repair and Prepare
||||| Strengthening Europe

Friedrichstraße 194
D – 10117 Berlin
Tel.: +49 (0)30/259 219 107

Online: delorscentre.eu
E-Mail: info@delorscentre.eu
Twitter: @delorsberlin

6/6