Artificial intelligence in Europe: balancing innovation with regulation

Sampsa Samila in conversation with Gabriele Mazzini

August 1, 2024

The EU AI Act, the world's first such legislation, was conceived to ensure safe,

transparent, traceable, non-discriminatory AI innovation. Will it succeed? The chief architect explains the thinking behind the new Act and its implications for the future.

After years of political wrangling, in May 2024, the European legislature officially adopted the Artificial Intelligence Act (EU AI Act), whose obligations will be rolled out across Member States over the coming months and years, coming into full effect by mid-2027.

This landmark legislation represents the world's first to regulate the fast-emerging field of AI. As such, like the EU General Data Protection Regulation (GDPR) before it, the EU AI Act may set the standards for other countries to follow. And as it is legally binding, with fines for non-compliance, there are obvious consequences for companies operating across borders, as we have already witnessed with the EU repeatedly fining U.S. tech giants for falling foul of its rules on a host of matters, from privacy to tax to competition.

Given the global ramifications, Sampsa Samila, director of the Artificial Intelligence and the Future of Management Initiative at IESE, was keen to discuss the Act with its chief architect and lead author, Gabriele Mazzini.

Mazzini isn't your typical career EU bureaucrat. Though he started working for the EU in 2001, he took a break in 2009 to pursue a master's in the U.S. and subsequently ended up working there for the Millennium Villages Project, an international development initiative using tech innovations to help lift low-income communities out of poverty. "That's where I developed an interest in technology and policy," he says.

He returned to Europe in 2017 wanting to work on the policy implications of emerging technologies. As AI began to take off, he was well positioned to provide expertise to the European Commission on the intersection of AI and EU law, and to lead the work on drafting AI regulation. He credits his experience of living in the U.S. and working in tech startups there with giving him an entrepreneurial mindset not commonly found in EU institutions.

In this interview conducted in the summer of 2024 as the Act began to go into effect, Samila and Mazzini discuss the ins and outs of the regulation — especially the delicate balance that must be struck between controlling the risks posed by AI while at the same time not excessively constraining innovation.

Goals of the EU AI Act

Sampsa Samila: What motivated the Commission to work on AI regulation? What were the key objectives?

Gabriele Mazzini: The Commission wanted to make sure we had a harmonized set of rules for AI-based products and systems across the entire EU, so their deployment and use in the EU would be compliant with EU law. A central tenet of the EU is the free circulation of goods and services throughout the internal market. And we need to ensure those goods and services meet minimum standards of safety and don't lead to undesired impacts (such as illegal discrimination) for there to be trust between companies and citizens. As such, we needed a common legal framework for there not to be barriers between different Member States and to provide legal certainty for businesses as they develop AI systems.

SS: How specifically does the Act protect these rights and principles?

GM: We established three categories of risk. First, if an AI system poses unacceptable risks, the law essentially foresees a prohibition. These would include applications involving cognitive behavioral manipulation, such as a voice-activated toy that encourages bad behavior in children, or social scoring that categorizes people, not on established practices like credit scoring, but on their "trustworthiness," for example.

Second, there may be high-risk AI systems, involving employee hiring and management, credit scoring and law enforcement, to name just a few, which risk violating people's rights. These are permitted but subject to certain compliance requirements, ex ante, and a procedure for conformity assessment. Most provisions of the Act fall under this category, and this imposes the greatest burden on companies.

The third category applies to AI systems posing a risk linked to a lack of transparency or disclosure — if humans aren't informed that they're interacting with AI while using a chatbot, for instance. This raises questions around human dignity and agency, which is another fundamental EU right.

EU's AI development strategy

SS: How does this play into the bigger picture of the EU's development of AI?

GM: The EU's AI strategy relies on two components: an ecosystem of trust, which is what we

have just been talking about, and an ecosystem of excellence, which primarily comes down to funding and coordinating R&D among Member States. It means investing in skills, innovation hubs and testing facilities for companies through programs like Horizon Europe, which complements the trust component. The two work together.

SS: How did you try to enact regulation without constraining innovation?

GM: This is the crux of any regulation. I think regulation can support innovation to the extent that it gives guidance. At the same time, any rule is a constraint. I wouldn't say that regulation always supports innovation, but I do believe it's a price you may have to pay for certain types of innovation.

We need to have this conversation. And different jurisdictions have different approaches. The U.S. is known for letting innovation go faster, in a freer form, and then relying on ex post market-driven remedies that either citizens or companies can activate through the court system — liability, for instance. In Europe, we tend to be more risk averse, to the extent that we want to prevent some possible negative externalities and consequences from the very beginning, when potential risks are just emerging. The potential cost of this approach is to constrain innovation by setting boundaries and safeguards as that innovation process is taking place.

The Act tries to strike this balance by relying on standards, rather than entering into technical details. By requiring companies to ensure compliance around data governance, documentation, human oversight, transparency and robustness, the Act remains fairly high level. This is a way of allowing the tech community to keep developing solutions, but that meet those requirements — supporting as much AI innovation as possible within a principled, stable legal framework.

AI regulation and smaller companies

SS: One big concern is that it's only the larger companies, mostly outside Europe, which have the resources to handle this kind of regulation and the costs that come with it. How do you facilitate the growth of smaller European companies with fewer resources?

GM: Personally, I'm not sure we managed to find the proper balance there. This would have required treating smaller players differently from big players. I recall some discussions

around whether there should be special treatment for SMEs to reduce the financial and other burdens on them. And certain ad hoc, albeit limited, provisions were introduced. But generally, I don't think the approach was bold enough, such as exempting certain market operators from some of the rules until they had reached a certain size. Member States and the Commission will need to undertake a number of activities in terms of training and providing dedicated communication channels to facilitate the participation of smaller actors in the standardization process. One thing smaller operators do have is priority access to the sandboxes.

SS: Can you tell us more about the role of the sandbox in promoting innovation?

GM: The sandbox provides a general framework for companies to develop and test AI systems in a controlled environment, meaning with the oversight of the authorities. This essentially offers a safer space to innovate while avoiding certain problems, like fines, if the company ultimately violates regulations while following the advice received. Companies will also have the possibility to test in real-world conditions (i.e., in a free environment) but with certain safeguards.

Regulations for foundation models

SS: How does the Act handle general-purpose or foundation models, which are in development and whose applications and capabilities we don't yet fully understand?

GM: This generated a lot of debate, especially during the final negotiations. In fact, the initial proposal didn't include any rules around foundation models. But with the emergence of ChatGPT, the European Parliament and Council felt this new phenomenon had to be addressed and introduced additional rules, although with a different focus and approach. The end result is a two-tier system, with rules that apply to all general-purpose models and additional rules that apply to those models that present a systemic risk. A critical question we had to answer very quickly was how to distinguish between these two. In the end, it came down to the amount of computing power used for their training as the only somewhat clear criterion to distinguish the two types of models.

SS: Why use compute rather than the size of the model or the size of the training data?

GM: The Biden Administration's Executive Order on AI in October 2023 put emphasis on measuring floating point operations (aka FLOPS), which relates to the compute needed for training the AI models. That played a key role in our thinking. Our understanding was that the FLOPS can be considered a sort of combined parameter that can account for other factors such as the amount of training data and the model parameters.

SS: As we get increasingly powerful models, is there a mechanism for updating this compute benchmark?

GM: Yes, there is. That was a crucial element to be added, to make this regulation future-proof. The first attempt to quantify FLOPS in the law can be modified. The Commission will have the power to change that threshold. What was considered a large model two years ago is certainly not considered large today, and going forward the parameters will change again.

It's also important to note that the AI Act foresees that, regardless of the compute threshold, the Commission may still have the power to designate certain models as posing a systemic risk based on a separate procedure, which considers other factors identified in an annex.

SS: Why did open-source models end up being excluded from these general-purpose AI regulations?

GM: They're not excluded, just covered less stringently. There are four sets of rules to consider here around: technical documentation that makes information available to authorities; transparency with downstream providers; and then two sets of obligations around copyright. These four obligations apply to all general-purpose AI models, with the exception that open-source models are exempt from two of them: technical documentation and transparency. Then, when it comes to models with systemic risks, there are additional obligations around risk assessment and risk management, which apply to all models regardless of whether they're proprietary or open-source.

Global influence of AI giants

SS: Let's talk about enforcement and global influence. Some, including Yann LeCun, have voiced concerns that, even if AI models are safe, there's a danger of them being run by a handful of giant companies. These companies would then have control over a huge chunk of our digital lives as well as the digital content we see.

GM: I agree with his point of view. We don't want to live in a world where our preferences are culturally influenced by a handful of people and organizations. We want to have a variety of digital content and preferences based on distinct cultures. And to attain this fair and varied digital diet, this would support the argument for open-source rather than the closed systems of just a few companies.

The EU motto is "united in diversity." We need to develop our own models that serve our own languages, our own cultural diversity. Why would we expect our own cultural diversity to be better served by a company based outside the EU? Foundation models, and language models in particular, represent a nascent market where openness is crucial and which provides a competitive advantage, first of all, for us in the EU, and more generally, for non-native English speakers. We have to be mindful, in that sense, of not closing the system, but having local actors that support and serve a certain language market and respect cultural nuances.

SS: It seems the risk-based concerns about security are independent of concerns about cultural and political influence. Does the regulation help facilitate the latter type of development?

GM: I think it does, though I do wonder if this specificity, and the value that Europe can add in that space, was sufficiently considered. That's not to say that the AI Act will be an obstacle to this type of development. Yes, the most stringent, top-tier rules apply to models with systemic risk. Nonetheless, we also have obligations for the lower tier. I hope that, over the next couple of years, the need to take into account cultural nuances will grow, and that will ensure the competitiveness of EU players developing models in native EU languages.

EU as groundbreaker in AI standards

SS: Apart from the Biden Administration's Executive Order, the European Commission is arguably the first government body to actually regulate AI. Is there more of a cost or a benefit to being the first to set standards that others may follow?

GM: I think it gives us a first-mover advantage, in the sense that we put a stake in the ground and, therefore, force everyone around the world to take that into account. Europe is a big market, and any company that wants to operate in it must abide by these standards. This also compels other jurisdictions — whether they have companies that operate in the EU or they simply feel the need to confront AI regulation for themselves — to look first at what is

happening here, and then decide how to relate to EU law as opposed to starting from scratch.

One important element to consider here is that we don't discriminate between EU companies and non-EU companies, because what matters when you want to sell your high-risk product in the EU market is that the same rules apply. Having said that, it's true that an EU firm will immediately be confronted with the AI Act and more stringent regulations in its domestic market, whereas a non-EU company may have more leeway in developing its products outside of the EU without the application of the AI Act and be able to grow before entering the EU. Companies confront those rules at different moments of their growth. And that brings us back to the matter of bigger companies having the means to comply with the regulation more easily.

SS: United in diversity, as you said, is the EU motto. How do we ensure consistent implementation and compliance across such diverse countries?

GM: In EU law, we, in Brussels, develop the harmonized laws, but the primary actors that are supposed to enforce them, with some exceptions, are the Member States. With the AI Act, we have a bit of a mixed situation. On the one hand, we have the national authorities of the Member States. There may be a risk of divergent application because some authorities may be more or less rigid. This risk is mitigated by the existence of a coordination mechanism, based on the EU Market Surveillance Regulation, which foresees that all market surveillance authorities meet on a regular basis and exchange best practices. Typically what they do is align their enforcement practices. In cases where there is conflict, the Commission can step in and decide for everybody.

SS: How do you intend to keep this regulation alive and relevant as AI technologies continue to develop?

GM: This was one of the concerns we had from the very beginning, back in 2019-20. The AI Act regulates a set of quite complex technologies and, as you said, AI has been evolving over time. We don't have a consensus, even among experts, about exactly what is and isn't AI. How do you take that complexity and build a legal framework, which by its nature requires a black-and-white approach? It must be clear which kinds of technologies the regulation applies to.

It's a difficult trade-off to, on the one hand, provide legal certainty for operators today and, on the other hand, enable flexible adaptation. We have to do this while remaining within a

stable and legally sound framework under our constitutional legal order in the EU. This requires that the essential elements of the law have to be established by the legislature and not the executive, which typically is more agile in its rule-making authority.

It’s an inherent difficulty. The technology moves so fast — faster than legislative actions or revisions. Remember that the AI Act took three years to be negotiated. Having said that, maybe the law can evolve a bit faster by changing some of the mechanics, some of the structural foundations. I think this will remain, in some respects, an unsolved issue. But we certainly will do the best we can.

This interview is based on two separate conversations conducted on June 21 and July 1, 2024, and edited for length and clarity.

This interview is published in IESE Business School Insight magazine #168 (Sept.-Dec. 2024).

www.iese.edu/insight