

The EU AI Act Explained: A Conversation with Professor Christoph Winter

October 2026



Introduction

The Simon Institute for Longterm Governance (SI) is publishing periodic conversations with leading legal and policy experts to explore how different countries and regions are navigating the governance of AI, and how emerging domestic practices might inform international AI cooperation. Our [first conversation](#) was with Professor Zhang Linghan on China's approach to AI governance.

In the second installment of our series, SI spoke with Professor Christoph Winter about the European Union's AI Act and the General-Purpose AI Code of Practice, and what they mean for AI providers outside the EU. Winter is Assistant Professor of Law and AI at the University of Cambridge, where he directs the Cambridge Programme on AI Science & Policy (CASP). He also co-founded and directs the Institute for Law & AI (LawAI).

He served as legal advisor to the Chairs and the European AI Office on drafting the General-Purpose AI Code of Practice, which sets out how providers of the most advanced AI models can comply with the AI Act.

Key takeaways:

1. **Enforcement of the AI Act has started, with stronger interventions expected.** The Commission has sent requests for information to more than 30 providers and received serious incident reports following recent incidents in which AI models reportedly escaped their testing environments. Winter thinks that, depending on how providers respond, fines are likely and requirements to withdraw models from the EU market could follow in some cases.
2. **The Act applies to providers placing general-purpose AI models on the EU market, wherever they are based.** Providers outside the EU must generally appoint an authorised representative in the Union. Providers of the most capable models face additional obligations to assess and mitigate systemic risks and report serious incidents. Whether these obligations cover internally deployed models is legally unsettled, but Winter advises providers to assume they do.
3. **The Code of Practice is technically voluntary but serves as the key benchmark for compliance.** Although providers can demonstrate compliance in other ways, non-signatories may nevertheless be expected to show how their measures compare with the Code. With harmonised standards some years away, Winter expects the Code to remain the main compliance route for the coming years, and notes it can evolve in line with technical developments.
4. **Limited regulatory capacity and political pressure could test enforcement.** Winter argues that, despite recent hiring efforts, the AI Office still needs significantly more technical and legal expertise. Enforcement could face resistance from foreign governments seeking to protect their companies, though Winter points out that the

obligations fall on companies, many of which have strong incentives to keep offering their models in the EU.

5. **The AI Act is only one component of European AI sovereignty.** Winter sees the Act as a robust foundation for addressing frontier AI risks, but argues that Europe also needs reliable access to advanced AI capabilities, defensive tools, and more resilient infrastructure. The EU's delayed access to Anthropic's Mythos, with its cybersecurity agency only gaining access to an older version months after launch, illustrates the challenges of relying on foreign providers to assess and respond to emerging threats.

This interview has been edited for length and clarity. The views expressed are those of the interviewee and do not reflect the official policy or position of SI. This publication is provided for general information only and does not constitute legal advice.

Contents

EU AI Act origins	3
1. Trajectory of EU digital regulation	
2. Why the EU regulated, and how the Act evolved	
The Act's design	4
3. Obligations for general-purpose AI providers	
4. Strengths and weaknesses	
5. Compute thresholds	
6. The General-Purpose AI Code of Practice	
7. Enforcement	
International implications and context	10
8. The Act's reach beyond the EU	
9. Advanced cyber capabilities	
10. Scope for cooperation on frontier risks	
Future directions	12
11. Implementation and the AI Office's constraints	

EU AI Act origins

1. How would you describe the trajectory of the EU's approach to digital regulation, and how does the AI Act fit within that tradition?

Over the past decade, the EU's approach to digital regulation has slowly moved away from decentralised enforcement towards centralised enforcement at Union level.

Earlier instruments, such as the [General Data Protection Regulation](#) (GDPR), left enforcement largely to the authorities of the Member States. When the [Digital Markets Act](#) (DMA) and the [Digital Services Act](#) (DSA) were adopted in 2022, the shift to centralised enforcement was clear. The DSA is especially noteworthy because it distinguishes between providers of online intermediary services in general and providers of very large online platforms or search engines, providing for centralised enforcement at Union level only in relation to the latter. The DSA also adopted the concept of 'systemic risk'; before then, it was commonly used in financial regulations. So the DSA marked an express recognition that if the Union wishes to ensure that systemic risks arising out of certain digital services are to be effectively constrained at Union level, enforcement must also take place at Union level.

The [EU AI Act](#) continues this approach with its rules on AI systems being largely enforced by Member States' national authorities, whilst its provisions concerning general-purpose AI (GPAI) models, including those with systemic risk, are enforced centrally at Union level. Put simply, under the AI Act, GPAI models are like foundation models: they have broad capabilities useful for a wide range of tasks. GPAI models with systemic risk can be understood as the most advanced, or frontier models; in other words, they are the most capable or powerful models, which present large-scale risks.

This approach of regulating the most powerful AI models at the Union level is further implemented in the AI Act by its establishment of the [European AI Office](#) within the European Commission (hereafter "the AI Office"). The AI Office is tasked, among other things, with the monitoring, supervision, investigation, and enforcement of the provisions concerning GPAI models.

2. Why did the EU decide to legislate on AI, and how did the Act evolve during the drafting process?

Two main considerations seem to have informed the initial push for EU-level regulation. First, the opacity, complexity, and autonomy inherent in AI systems meant that existing rules would be difficult to apply directly, particularly where AI puts fundamental rights and public safety at risk. Second, several Member States were considering adopting national rules to regulate AI. Had they done so, this would very likely have led to fragmented and potentially incompatible regulatory approaches across the Union, which meant that harmonisation at Union level was highly desirable.

In terms of how that framework took shape, one of the most significant developments in the AI Act's drafting process was the relatively late inclusion of GPAI models and GPAI models with systemic risk.

Notably, the European Commission's initial proposal for the AI Act in 2021 dealt only with 'narrow' AI systems. However, during the legislative process, the co-legislators (the European Parliament and the Council) recognised how important it was or would be to regulate AI at the level of GPAI models and their providers. The Council first proposed the concept of a 'general-purpose AI system' in late December 2022, shortly after ChatGPT was first launched. That proposal was then supplemented by the European Parliament's proposal introducing the idea of a 'foundation model'. During the legislative debates between the three institutions in late 2023, these concepts then evolved into the GPAI model terminology we see in the AI Act today. The inclusion of GPAI models has directly contributed to the Act being the first and, as of today, possibly still the best-placed regulatory instrument in the western world to respond to the kinds of risks and incidents we are increasingly seeing materialise from frontier AI.

The overall perception of risks from AI also shifted markedly during the legislative process. The public debate following the release of ChatGPT intensified lawmakers' focus on general-purpose AI and made certain systemic risks especially salient among the legislators. This is clearly visible in some of the recitals of the adopted AI Act.¹ Indeed, Recital 110 specifically refers to chemical, biological, radiological, and nuclear risks, as well as offensive cyber capabilities, that were absent from the European Commission's original proposal.

The resulting Chapter V of the AI Act, and especially Article 55 on the obligations of providers of GPAI models with systemic risk, has provided a robust legislative foundation that requires providers to take systemic risks seriously and to develop appropriate safety and security frameworks in line with the state of the art. The use of codes of practice under Article 56 then provided a mechanism to specify how companies can comply with the often rather broad requirements of the GPAI provisions. Importantly, it also ensures that the EU isn't locked into a rigid legal framework, unable to respond to rapid developments within AI.

The Act's design: GPAI and the Code of Practice

3. For GPAI developers who want to access the EU market, what are the obligations they most need to understand?

The AI Act's GPAI obligations have been active since August 2025, and enforcement took effect from 2 August 2026. Crucially, all GPAI models, regardless of where their providers are based or where the models are developed, fall within the AI Act's scope once they are placed on the

¹ Editor's note: In EU law, recitals form the non-binding preamble to a legal act; they explain the context, legislative intent, and objectives behind its articles.

Union market.² The most important provisions for providers of GPAI models are Articles 53–55 of the AI Act.

Article 53 sets out the obligations that apply to all GPAI models. It covers the technical documentation that providers must prepare for the purposes of the AI Office’s supervision, including detailed information on the size and capabilities of their models and on the evaluations carried out on them. It also covers the documentation that providers must make available to downstream providers who integrate their models into their own AI systems, as well as the copyright obligations providers must follow.

In addition to Article 53, Article 55 applies specifically to providers of models classified as GPAI models with systemic risk, which again, is EU-speak for “frontier AI” and “frontier AI risks”. It stipulates that providers of these models must commission and/or carry out evaluations in order to assess and mitigate systemic risks appropriately. Article 55 also requires providers to record serious incidents involving their models and report them to the AI Office.

[Recent incidents](#), in which AI models reportedly escaped their testing environments and gained unauthorised access to outside systems, have demonstrated how important such disclosures are for the AI Office’s supervisory functions. This is not just because they may reveal non-compliance with the AI Act’s obligations, but also because they allow the AI Office to build a better understanding of the risks these models present in real-world scenarios.

Article 54 is also relevant for GPAI model providers established outside the Union. Before placing a GPAI model on the Union market, those providers must ‘by written mandate’ appoint an [authorised representative](#) established in the Union. That mandate should empower the representative to carry out certain compliance-related tasks enumerated under Article 54(3).

The EU team at LawAI, together with the Cambridge Programme on AI Science & Policy (CASP), was instrumental in putting together the Cambridge Commentary on EU General-Purpose AI Law. The Commentary examines these provisions, together with the other AI Act rules relevant to GPAI models and their providers, with deep legal rigour, and offers a balanced assessment of possible interpretations of the AI Act’s rules. The Cambridge Commentary is freely available online at <https://cambridge-commentary.ai> and is periodically updated to reflect relevant legal developments. New chapters are also being published on an ongoing basis.

How do the AI Act’s obligations apply differently to open-source models and what is the rationale for this approach?

GPAI models released under a free and open-source licence are exempt from the documentation obligations under Article 53(1)(a)–(b), which cover the technical documentation for the AI Office and the information for downstream providers, as well as the obligation to

² Editor’s note: Under Article 3(9) of the AI Act, “placing on the market” means the first making available of an AI system or general-purpose AI model on the Union market. “Making available” (Article 3(10)) covers supply for distribution or use in the course of a commercial activity, whether in return for payment or free of charge.

appoint an authorised representative under Article 54. However, these exemptions don't apply if the open-source model is classified as a GPAI model with systemic risk.

The exemptions depend on the licence and the information made available, rather than simply on whether a provider describes its model as "open". To benefit from them, the GPAI model's licence must permit "access, usage, modification, and distribution of the model". Its "parameters, including the weights, the information on the model architecture, and the information on model usage" must also be made publicly available.

The rationale is that the open release already provides a degree of transparency that reduces the need for those particular documentation requirements. The exemption does not, however, extend to the copyright policy and the public summary of training content, which apply to all providers.

For models with systemic risk, the concern goes beyond access to information. Providers must evaluate their models, assess and mitigate systemic risks, report serious incidents, and ensure adequate cybersecurity. Making a model openly available does not establish that these obligations have been met. The Act therefore recognises the benefits of openness without treating it as sufficient to address the risks posed by the most capable models.

4. What do you see as the key strengths and weaknesses of the AI Act?

Whatever you think of other parts of the AI Act, its chapter on general-purpose AI and the [GPAI Code of Practice](#) are the world's first frontier AI rules, and I believe these parts are well done, though I should disclose that I was also the legal advisor to the Chairs and the AI Office for the drafting of the Code of Practice.

A particular strength of the current version of the AI Act, and especially of the framing of its obligations under Article 55 for GPAI models with systemic risk, is its adaptiveness regarding systemic risk assessment and mitigation (see question 6). This allows the Act to naturally evolve as the models themselves evolve and as techniques for better identifying, assessing, and mitigating systemic risks develop.

A weakness that has been discussed in the past, but has become especially salient in recent months, is the question of how internally deployed models are governed under the AI Act. The AI Act was meant to be risk-based, and arguably, some of the largest risks are being posed by internal deployment, but whether internal deployment is actually covered by the Act is legally uncertain.

The answer to this is likely to become clearer as the European Commission's practice develops. But in the meantime, it seems to me that it is in the providers' own interest to assume that internal deployment is in scope of the AI Act given its emphasis on risk assessment and mitigation throughout the lifecycle of a GPAI model. This would also ensure that they would not be in breach of their obligations if those models come within the AI Act's scope at a later point.

5. Compute thresholds are central to the AI Act. Why were these specific thresholds chosen and how do they work in practice?

Compute thresholds are useful proxies for the capabilities and risks of AI models. This is because training compute is relatively easy to measure and it has tended to correlate with emerging capabilities so far.

The most important threshold is 10^{25} floating-point operations (FLOP). Under Article 51(2) of the AI Act, models trained using a cumulative amount of compute greater than 10^{25} FLOP are presumed to have high-impact capabilities. This means that these models are classified as GPAI models with systemic risk under Article 51(1). The figure was chosen because it was thought to roughly mark the point at which models develop the kinds of capabilities that could give rise to systemic risks. The Annex to the Commission's [GPAI Guidelines](#) sets out in more detail what the Commission counts towards cumulative compute and how it should be calculated.

It is worth stressing that this threshold is a presumption, and it works in both directions. A provider whose model crosses 10^{25} FLOP can argue that the model does not in fact have high-impact capabilities. Equally, under Article 51(1), a model trained with less compute can still be classified as a GPAI model with systemic risk if it does in fact have such capabilities. The European Commission can update this threshold as the technology develops.

There are two other compute thresholds worth discussing. They are used by the European Commission as so-called indicative criteria in its GPAI Guidelines.

The first, 10^{23} FLOP, helps determine whether a model counts as general-purpose at all. Models above this level are likely to qualify, although the Commission also looks at other factors, such as the kinds of inputs and outputs the model works with (text, images, audio) and the range of tasks it can perform.

The second compute threshold deals with a question that matters a lot in practice: when does someone who modifies an existing model, for example by fine-tuning it, become the provider of a new model, with all the obligations that come with that? The European Commission in this case created relative thresholds that are based on how much compute was used to modify the model, compared to the amount of compute originally used to train it. As a rule of thumb, if the modification uses more than a third of the compute originally used to train the model, the result is treated as a new model. The idea is that a modification of that scale could plausibly give the model new capabilities, and potentially new systemic risks.

Again, this is an indication rather than a strict requirement. If someone fine-tunes an existing model with a very small amount of compute, but nevertheless produces very significant changes in the model's capabilities or systemic risks, the European Commission may still consider the result to be a new model.

6. How should non-EU providers understand the legal status of the General-Purpose AI Code of Practice (hereafter “the Code”)? How does it function alongside the Act, and what does signing entail in practice?

As I mentioned earlier, the AI Act deliberately relies on co-regulatory instruments³ to set out more specifically how the Act’s obligations should be complied with in practice.

The Code’s legal status is rather complicated. It is technically a voluntary instrument, but its effects in practice are unusually strong for a few reasons.

First, unlike harmonised standards, which give providers what lawyers call a “presumption of conformity”, following the Code doesn’t, on its own, guarantee that a provider complies with the AI Act. However, the European Commission and the AI Board (made up of Member States’ representatives) have communicated that they consider the Code “adequate” in covering the obligations under Articles 53 and 55. So providers have strong assurance that the AI Office will treat compliance with the Code as a way of demonstrating compliance with the relevant provisions of the AI Act.

Providers don’t have to sign the Code. Instead, they can show compliance through what the Act calls “alternative adequate means”. But the Code’s comprehensiveness, and the fact that it was drafted with input from industry, other stakeholders, and experts on how best to put the Act’s obligations into practice, make it unsurprising that many of the leading frontier AI companies have signed it.

Second, the European Commission has said that non-signatories may be expected to explain how their own measures ensure compliance, for example through a gap analysis comparing them with the measures in the Code. This means the Code is used as a benchmark even if a GPAI model provider hasn’t signed it. Signing also brings a practical benefit: the AI Office offered closer cooperation to providers adhering to the Code during its first year of implementation, helping them work towards full compliance without delays in placing their models on the Union market.

How will the Code’s role evolve if harmonised standards arrive?

Articles 53(4) and 55(2) of the Act provide that providers may rely on codes of practice to demonstrate compliance “until a harmonised standard is published”. Compliance with harmonised standards, by contrast, gives providers a presumption of conformity, at least to the extent that those standards cover the relevant obligations. Once such standards are published, they will therefore likely take precedence over codes of practice as the main route to demonstrating compliance. Alternative adequate means of compliance will, however, remain available even once harmonised standards exist.

³ Editor’s note: Co-regulatory instruments, such as the General-Purpose AI Code of Practice, are governance mechanisms developed with input from industry and other stakeholders, and endorsed by public authorities.

For now, however, this is a somewhat theoretical question. Harmonised standards for GPAI models still seem a long way off, and for the foreseeable future, I believe the Code will remain the main instrument on which providers can rely to guide and structure their compliance efforts.

What mechanisms are there to allow the Code to evolve if necessary?

There are two main mechanisms through which the Code can evolve. The first is formal review, in which the AI Office plays the central role. Under Article 56(8), the AI Office is tasked with encouraging and facilitating the review and adaptation of codes of practice, “in particular in light of emerging standards”. An amendment is very much a possibility if the current Code comes to be seen as no longer reflecting an accurate understanding of how Articles 53 and 55 should be operationalised, especially in light of technical developments.

For this, the Chairs and Vice-Chairs of the Code’s Safety and Security Chapter have recommended that the Commission set out a mechanism for updating the Code, with reviews on a regular cycle, such as every two years, to give providers predictability. They have also suggested that emergency updates with immediate effect should be possible, though only in very rare and exceptional cases. These are good suggestions.

The second mechanism is built into the Code itself. The Code has been drafted to allow considerable flexibility to capture new developments without needing to be amended, such as via the inclusion of dynamic references. The various requirements that refer to the “state of the art” are a good example of such dynamic references. The AI Office can also issue guidance on how to interpret these references as the technology and risk mitigation mechanisms evolve.

7. What enforcement actions have we seen already and what scenarios could you see playing out over the coming 6–12 months?

Since the AI Act’s enforcement powers took effect on 2 August 2026, the Commission has started to use them. It [recently confirmed](#) that it had sent requests for information to more than 30 providers. Following recent incidents, the AI Office has also requested and received serious incident reports. It is now preparing further requests to investigate these incidents, which the [AI Act Scientific Panel](#) will also inform.

It is very difficult to predict the next six to twelve months, in part because we don’t know how many and how severe the incidents will be, and what the political reactions will be. The default is that the European Commission will escalate its enforcement.

Depending on the reactions of the providers in question, the Commission might require them to take their models off the market and/or issue fines. Given the current pace of development and lack of safety mechanisms in many places, I think this is very likely to happen in at least some instances. While there are political dynamics that make taking some of these measures more complicated, the Commission will lose credibility if it doesn’t enforce appropriately when models pose serious risks to the European people.

International implications and context

8. What is the legal basis for the AI Act's application to providers outside the EU?

Under Article 2(1)(a), the Act applies to GPAI models placed on the Union market, irrespective of whether providers are established in the Union or elsewhere. A company does not have to be European for the EU to set conditions on the products it offers to people in Europe. This is a common approach to product regulation in the EU and other jurisdictions.

I think it's important to distinguish between applying rules to foreign providers and regulating conduct outside the Union. The former is not unusual. But compliance with those rules can, of course, affect the latter. A company developing a model in California may need to change aspects of its development if it wants to offer that model in Europe.

Do you think governments outside the EU might resist enforcement actions taken by the AI Office?

It is possible, particularly if enforcement would require a leading provider to take a model off the Union market. The provider's government might view that as a measure against a strategically important domestic company and put pressure on the Commission to reconsider.

But government resistance will not necessarily prevent enforcement. The AI Act's obligations fall on companies, many of which have strong incentives to continue offering their models in Europe and to comply with the conditions for doing so. A provider may therefore seek to resolve the Commission's concerns even while its government objects.

How much political pressure would constrain the Commission is difficult to predict, and will heavily depend on whether the Commission believes the model poses serious risks. If the Commission actually takes Loss of Control or Chemical, Biological, Radiological, and Nuclear (CBRN) risks seriously, it is much less likely to be pressured into holding off on meaningful enforcement that could protect the European people.

9. Frontier models such as Anthropic's Mythos have demonstrated advanced cyber capabilities. How has the EU responded, and how prepared is the regulatory regime for models with these capabilities?

The EU has taken meaningful steps to obtain access and information, but I would be careful about treating those steps as evidence that it is adequately prepared. In September, the Commission confirmed that the EU Agency for Cybersecurity (ENISA) had access to Mythos 5 and was testing it. It also said that it had begun using requests for information to check providers' compliance. Those are useful and important steps, but the important question is what the institutions can do with that access and information.

The Act and the Code already provide a basis for addressing systemic risks arising from advanced cyber capabilities in a reasonable way. The bigger challenge, however, is applying them effectively, especially as capabilities accelerate.

The EU certainly faces many challenges when it comes to advanced cyber capabilities. Many relevant public institutions and companies were not initially included in [Project Glasswing](#), Anthropic's programme that gave some actors early access to Mythos to patch their cyber vulnerabilities. ENISA only gained access several months after launch, and not to the latest version of the model. Without access to such programmes, EU institutions may be left vulnerable as cyber capabilities proliferate.

This raises broader issues of European sovereignty and of dependence on foreign technologies. The AI Act is an important part of the response, but it needs to be complemented by reliable access to defensive tools, technical expertise, and more resilient infrastructure. The recent report '[A Transformative AI Strategy for Europe](#)', to which I and other LawAI colleagues contributed, addresses precisely these needs. It recommends using Europe's position in the AI supply chain to secure access to frontier capabilities, bringing AI expertise into public institutions, and preparing critical infrastructure for AI-enabled attacks. These measures would strengthen Europe's ability both to assess emerging threats and to respond to them.

10. There is growing international interest in frontier risks such as CBRN misuse by non-state actors. Could this support cooperation on monitoring or incident reporting?

There is certainly growing interest in responding to frontier or systemic risks across the world, and this interest creates good opportunities for cooperation. For example, the AI Office already draws on independent third-party evaluations from institutes such as the UK AI Security Institute to inform its own monitoring activity. And the Code requires signatories to include any available reports from independent external evaluators in the Model Reports they provide to the AI Office.

Interestingly, the Act does provide an explicit route for certain exchanges of confidential information with regulators outside the EU. Article 78(5) allows the Commission and Member States to exchange such information where necessary under certain conditions. So there is a legal basis for exploring practical cooperation. That said, I would be cautious about assuming that shared concerns will definitely translate into extensive information-sharing, particularly where commercial or national security interests are involved.

Future directions

11. Looking ahead, what can we expect for the AI Act and Code of Practice over the next two years? What would successful implementation look like?

I would expect the practical meaning of the Act and the Code to become clearer as the AI Office assesses providers' compliance and responds to particular risks as they emerge. It is hard to predict which models or incidents will drive that process. The Act's flexibility is a strength, but that flexibility only helps if providers and regulators can recognise when their approach needs to change and act on that understanding.

I think implementation will be going well if:

- Risks are appropriately identified and mitigated early enough to prevent serious harm. This could be before models become widely available, and some of this may depend on whether the Commission interprets the GPAI systemic risk provisions to include internal deployment,
- Providers communicate responsibly, proactively, and transparently with the AI Office about the risks,
- When incidents do occur, they are quickly identified, reported, and taken into account in providers' safety and security frameworks, and
- The AI Office has sufficient capacity to assess that information and, where necessary, secure changes in providers' conduct.

What role will the AI Office's practical constraints play? Will it have the technical staff, budget, and model-access arrangements to run meaningful evaluations?

The current constraints on the AI Office could make a substantial difference to how effectively the Act is enforced. The AI Office needs to assess whether providers' evaluations are adequate, identify risks those evaluations may have missed, and judge whether the proposed mitigations are sufficient. To do this, and do it well, requires the AI Office to have considerable technical and legal expertise.

The AI Office is staffed by an excellent team and is continuing to recruit, in particular technical and legal staff who are willing to forgo careers in the private sector in order to work on these very important issues at the European Commission. Its recent announcement that it will hire an additional 40 staff members is a step in the right direction. But given the scale and importance of the task at hand, I don't think this will be sufficient, and substantially greater investment is needed.

Talent is very important, and governments currently compete with companies that can offer much higher compensation and generally more attractive working conditions. Resources for more competitive salaries need to be accompanied by faster recruitment procedures and better ways to draw on external expertise.

This is a central argument of [‘Radical Optionality’](#), which I co-authored with Charlie Bullock. We do not know exactly which actions governments will need to take over time, but we do know that developing the expertise to act competently in various scenarios takes time, and that it is urgent.

More generally, I think governments across the world should be investing substantially more in their capacity to understand and regulate frontier AI well. Given the scale of the potential benefits and risks, that investment is justified even if the most consequential capabilities take longer to emerge than one might expect based on the past few months. Governments should have been building this capacity yesterday, but today is still better than tomorrow.

