CEDPO

CONFEDERATION OF EUROPEAN DATA PROTECTION ORGANISATIONS

Bonn, Bucharest, Dublin, Lisbon, Madrid, Milan, Paris, The Hague, Vienna, Warsaw

What's a FLOP?
How the AI Act Regulates General Purpose AI Systems

CEDPO AI and Data Working Group
Micro-Insights Series
March 2024

Authors:
Maria Moloney
Jared Browne

Contact information:
https://cedpo.eu
info@cedpo.eu

About the Micro-Insights Series

The 'Micro-Insights' Series is a publishing initiative by the CEDPO AI and Data Working Group. It will provide digestible, definitive, short-form papers on key areas of interest at the intersection between data and artificial intelligence. With a practical focus, keeping one eye on explaining complex topics and the other on implementation, it will outline the significance of key areas and advise practitioners on impact, and next steps. With the EU Artificial Intelligence Act (the 'AI Act') coming into law in 2024, the scene is now set for all practitioners, and it is possible to discuss the regulation of data and AI with much greater clarity.

The Micro-Insights Series will follow the evolution of AI and data over the coming years, and as the clock winds down on the crucial implementation period for the AI Act, and as AI technologies evolve in ever-more novel and unexpected ways, the Series will respond with up-to-date, authoritative guidance on the core areas of concern.

Amongst others, the series will include papers on:

- Regulation of General-Purpose Artificial Intelligence under the AI Act
- Explaining the AI Pact.
- Educating practitioners on how to conduct Fundamental Rights Impact Assessments under the AI Act.
- Outlining the role that data protection regulators will have in AI regulation.
- Examining whether or not the data protection officer is the right person to be the AI officer.
- The lawful basis for using training data in machine learning.
- Readiness toolkit for the AI Act.

Table of Contents

1. Regulation of General Purpose AI under the AI Act

In the context of the AI Act, a General Purpose AI (GPAI) system is defined as an AI system that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model was originally released on the market. These AI systems can be integrated into a variety of downstream systems or applications. GPAI models are also often referred to as Foundation Models. They have purposefully been designed to perform a wide range of tasks and to easily adapt to new situations. To this end, they are trained on very broad sets of unlabelled data and can be used for many different tasks without much fine-tuning.

In simpler terms, GPAIs are AI systems that can be used in and adapted to a wide range of applications for which they were not originally, intentionally or specifically designed[1]. This includes AI models like ChatGPT, DALL-E and Gemini by Google. These systems can perform a variety of tasks like generalising and understanding, operating across, and combining different types of information.

GPAI models are being used more and more as a type of digital infrastructure by companies to provide their customers with downstream services. Models like GPT-3 and Gemini are used to power chatbots and virtual assistants, providing users with interactive and dynamic responses. Other GPAI models like Hootsuite can generate content for a variety of purposes, such as writing articles, creating advertisements, or generating social media posts. Companies like Alpha Cephei are using open-source AI models for speech recognition and developing them into enterprise software products[1]. In healthcare, GPAI models are helping medical experts to analyse medical images, predict disease progression, and personalize treatment plans for patients[1]. While in finance, these models are used for credit scoring, fraud detection, and algorithmic trading.

In addition to these applications, other GPAI models are increasingly being used in fields such as creative writing, art generation, computer coding, medicine and healthcare, the life sciences, and even in programming and machine learning fields. The flexibility and adaptability of GPAI models make them suitable for a vast range of applications, and we can expect to see them used in even more innovative ways in the future.

The AI Act imposes certain obligations on providers of GPAI systems, with an aim to regulate the development and use of AI by providing a set of obligations for parties involved across the entire AI supply chain. The Act proposes a tiered approach to regulating GPAI systems. The tiered approach aims to balance the need for regulation with the massive potential of AI, ensuring that the technology can be developed and used responsibly while mitigating any risks.

This approach is based on the level of risk posed by the GPAI models and systems, with more obligations for those GPAI model providers that entail systemic risks[2].

---

[1] General-purpose artificial intelligence (europa.eu)
[2] Regulating Foundation Models in the AI Act: From "High" to "Systemic" Risk - MIAI (ai-regulation.com)

2. What obligations does a 'systemic risk' designation create?

The tiered approach to regulating GPAI models means lower or higher compliance obligations depending on the power of the model involved. This creates a baseline compliance level for providers of GPAI models, with all such models automatically being considered high risk, while a specific subset will be considered to meet the higher threshold of 'systemic risk' and face a heavier compliance burden. The below table outlines the respective compliance obligations for these models:

Tiered Compliance Obligations of GPAI Models under the AI Act

Requirement	Tier 1: All GPAI	Tier 2: GPAI with Systemic Risk
Maintain technical documentation of model, including training and testing processes, along with results of evaluations.	✔️	✔️
Maintain documentation for other AI providers that plan to integrate with the GPAI, including detail of specification and limitations.	✔️	✔️
Establish a policy for complying with EU copyright law.	✔️	✔️
Publish a detailed summary of the content used for training the AI model.	✔️	✔️
Co-operate with the EU Commission and national authorities in the execution of their duties under the AI Act.	✔️	✔️
Perform model evaluation.	✖️	✔️
Assess and mitigate possible systemic risks.	✖️	✔️
Report serious incidents to relevant authorities.	✖️	✔️
Adopt corrective measures to address the risk associated with any identified vulnerabilities.	✖️	✔️
Ensure an appropriate level of cyber-security for the model and its physical infrastructure.	✖️	✔️

3. How is 'systemic risk' understood and calculated?

GPAI models that entail systemic risk are deemed to have high impact capabilities, such that any negative incidents could have disproportionate effects on the technology value chain of which they are a part, and on the businesses, organisations and end users that may come to rely on them.

How then is 'high impact capability' assessed? High impact is understood to apply to any model where the cumulative computing power used for its training is greater than 10^25 floating point operations, or FLOPS. FLOPS per second measure a computer's processing speed. The basic chain of logic here is that more FLOPS means more power which means more risk.

Developers can challenge this assessment, where though the model may meet the 10^25 threshold, but, in practice, there may be specific circumstances which show that systemic risk is not the correct assessment.

FLOPS are not a sole performance indicator and the EU Commission's AI office will be able to set additional benchmarks for assessing systemic risk. The EU Commission also has delegated powers to amend the FLOPS calculation as well as to create new benchmarks. For example, this may take place in response to technological advancements in algorithms or enhancements to hardware. Additionally, the EU Commission, in assessing what systemic risk means, could consider the number of end users; the number of registered businesses using the model in the EU; the modality of the model; and the number of tokens and parameters used in the model.[3]

4. Who will oversee 'Systemic Risk' GPAI Models?

The EU AI Office will be set up by the EU Commission to oversee the most powerful GPAI models. This office will also be charged with creating and disseminating standards and testing practices and enforcing the common rules in all Member States. The AI Office will be supported by an independent panel of scientific experts.

The AI Board, made up of Member State representatives, will also have a role as a co-ordinating and advisory body and will assist Member States with implementation of the rules. An AI advisory forum will also be formed to provide expert advice to the AI Board.

---

[3] Regulating Foundation Models in the AI Act: From "High" to "Systemic" Risk, Cornelia Kutterer. https://ai-regulation.com/regulating-foundation-models-in-the-ai-act-from-high-to-systemic-risk/

5. Do any models currently meet the threshold of 10^25 FLOPS?

In real technological terms, very few systems will meet 10^25 FLOPS and the accepted consensus is that perhaps only OpenAI's ChatGPT4 and Google's Gemini models would currently be in scope. Such a high threshold for applying the most onerous level of compliance obligations means that very powerful systems, such as ChatGPT3.5 will, to some extent, slip though the net and only face mainly technical documentation-type obligations, and will not, for instance, have to monitor, address and report major risks.

6. When will these rules apply?

Provider of GPAI models will have 12 months from the coming into law of the AI Act to comply with their obligations.