October 20, 2022

EU AI Act – GitHub Position Paper

Every day, AI researchers and software developers across Europe contribute to AI innovation. To do so, they publicly share AI-related code and artifacts, in order to collaborate with a diverse, distributed community of researchers, academics, individual developers, governments, and businesses of all sizes. Open source software plays an important role in these collaborations, as do software development and sharing platforms like GitHub. GitHub is the world's largest such platform, with more than 22 million developers in Europe and 83 million worldwide.

EU developers use platforms like GitHub to share open source code and collaborate on all layers of the AI software stack (see graphic detail below). It is common practice for scientific AI research to yield a research paper manuscript published on the open-access academic repository arXiv, with proof-of-concept code shared online via GitHub. It is also common practice for AI development-oriented code to be built and shared publicly online, including on GitHub. This type of open collaboration has been key for enabling EU developers, individually and through businesses, to cement Europe as a key player in AI research, development, deployment, and commercialization. It is in support of this community that we write this brief with recommendations to ensure the AI Act remains a driver for European AI innovation.

The AI Act promises to support AI innovation in the single market while addressing risks posed to fundamental rights and safety. As debates on the AI Act continue, it is essential that AI researchers and developers be protected in their activities insofar as they do not constitute providing an AI system per se. Improved clarity on research and development activities, especially as they pertain to general purpose AI systems, can help deliver the stated goals of the Act.

Improve the Act by increasing clarity on AI research and development

We offer five recommendations for changes to the Act text can more clearly protect AI research and development activities within the single market:

Recommendation 1:
Add Recital text that acknowledges AI-related code is not an AI system

Recital 60 of the original Commission proposal[1] provided useful clarification that "software, software tools and components, pre-trained models and data, [and] providers of network services" are not making an AI system available on the market per se. We recommend adding clarifying text in Recitals, so that actors that build, often in research or volunteer capacities, and

[1] COM(2021) 206 reads "In the light of the complexity of the artificial intelligence value chain, relevant third parties, notably the ones involved in the sale and the supply of software, software tools and components, pre-trained models and data, or providers of network services, should cooperate, as appropriate, with providers and users to enable their compliance with the obligations under this Regulation and with competent authorities established under this Regulation."

host AI-related code do not face uncertainty. This will help focus regulators on AI systems made available on the market as opposed to code related to a single part of the AI software stack. Suggested language below reflects the latest Council Presidency text with ***additions in italicized bold***.

- Recital 6: The notion of AI system should be clearly defined to ensure legal certainty, while providing the flexibility to accommodate future technological developments. The definition should be based on key functional characteristics of artificial intelligence distinguishing it from simpler software systems and programming approaches. In particular, for the purposes of this Regulation AI systems should have the ability, on the basis of machine and/or human-based data and inputs, to infer the way to achieve a given set of objectives using machine learning and/or logic- and knowledge based approaches and to produce outputs such as content for generative AI systems (e.g. text, video or images) predictions, recommendations, or decisions influencing the environment with which the system interacts, be it in a physical or digital dimension. ***Software code and AI model objects, including pre-trained models, that are incapable of producing outputs directly should not be considered AI systems within the meaning of this Regulation.*** A system that uses rules defined solely by natural persons to automatically execute operations ***also*** should not be considered an AI system. AI systems can be designed to operate with varying levels of autonomy and be used on a stand-alone basis or as a component of a product, irrespective of whether the system is physically integrated into the product (embedded) or serve the functionality of the product without being integrated therein (non-embedded). The concept of the autonomy of an AI system relates to the degree to which such a system functions without external influence.
- ***Recital 60: In the light of the complexity of the artificial intelligence value chain, relevant third parties, notably the ones involved in the sale and the commercial supply of software, software tools and components, pre-trained models and data, or providers of network services, should cooperate, as appropriate and proportionate to risk, with providers and users to enable their compliance with the obligations under this Regulation and with competent authorities established under this Regulation.***

The AI Software Stack

The AI Act focuses on AI systems[2] put into service in the single market. Several layers[3] of software are required in order for an AI system to be put into service. The AI-related code in these layers are individually out of scope of the Act. For example, the AI model artifact alone is insufficient to put an AI system into service.

Interface
Provides a way for the system to interact with the environment, whether by way of graphical user interface, command line, or another mode. Example: Streamlit

Model serving and monitoring
Enables the AI model artifact to operate. This may take the form of APIs enabling remote access to large AI models or local access for smaller ones. In both cases, data pipelines must be specified in order for the system to be able to operate. Example: TensorFlow Serving

AI model
Within the machine learning paradigm, the AI model is an artifact that results from training an algorithm on data. Depending on their size, these code artifacts may be stored on GitHub directly or on third-party cloud storage and other web hosting services. Example: GPT-NeoX

Training and evaluation
Numerous software packages provide resources to manage AI model training and evaluate resulting models. These include tools to monitor performance on responsible AI features like bias. Examples: Sacred, Fairlearn

Algorithm selection
At the outset of AI model training, the algorithm must be specified. This is commonly done in custom software code that makes calls to frameworks like TensorFlow or PyTorch. Example: StyleGan 2

Infrastructure management
In order to perform training, especially for large datasets, resource virtualization is required to distribute the computational load across multiple processors, whether locally or in the cloud. Example: Kubernetes

[2] As defined by the latest Council compromise draft 2021/0106(COD) 12549/22 in Article 3(1): an AI system is “a system that is designed to operate with a certain level of autonomy and that, based on machine and/or human-provided data and inputs, infers how to achieve a given set of objectives using machine learning and/or logic- and knowledge based approaches, and produces system-generated outputs such as content (generative AI systems), predictions, recommendations or decisions, influencing the environment with which the AI system interacts;”

[3] It does not portray the AI value chain, as each layer can be coded or run by the same or different actors. It also does not depict many tasks associated with data collection and cleaning. Data management is invoked in both training (4-5) and deployment (1-2). Training yields artifact 3.

Recommendation 2:
Clarify what constitutes a "distributor"[4]
Adding the Recital text from Recommendation 1 can help provide clarity that open source software development and sharing platforms, including GitHub, which make AI-related code and content available do not constitute distributors "that makes an AI system available on the Union market".[5] As acknowledged by the Commission's Product Liability Directive proposal 2022/0302 (COD), Recital 28 "... When online platforms play a mere intermediary role in the sale of products between traders and consumers, they are covered by a conditional liability exemption under the Digital Services Act." Similarly, Recital 12 "... source code of software, however, is not to be considered as a product for the purposes of this Directive as this is pure information." Just as logistics companies that play a mere intermediary role in conveying components of products between economic actors, so too do open source software development and sharing platforms convey AI-related code and content without distributing AI systems as such.

Recommendation 3:
Amend Article 2(7) to exempt code artifacts and sharing activities
Together with new Recital text referenced above, Article 2(7)[6] should expand to include "and resulting artifacts" in order to clarify that research and development activities involve the sharing of code and that the underlying code artifacts themselves are not subject to the regulation as fully formed AI systems.

- "This Regulation shall not apply to any research and development activity regarding AI systems***, including the sharing of free and open source software code and AI models***."

Recommendation 4:
Amend Title IA General Purpose Systems to reconcile with Article 2(7)
If Title IA remains, we request to strike Article 4a(2)[7] to preserve the research and development exemption for AI model artifacts that may be further fine-tuned before being put into service. Further fine-tuning invokes development of an AI system by modifying its model through exposure to additional training data. This is ultimately a development activity. As currently worded, the Act risks pulling AI researchers or individual developers into scope who may share an AI model artifact publicly, even if they do so without intention of becoming a commercial supplier or placing an AI system into service. As recently proposed by the Commission, the

[4] In the latest Council compromise draft 2021/0106(COD) 12549/22, Article 3(7&10) identifies a distributor as "any natural or legal person in the supply chain, other than the provider or the importer, that makes an AI system available on the Union market," meaning "any supply of an AI system for distribution or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge".
[5] Latest Council compromise draft 2021/0106(COD) 12549/22, Article 3(7).
[6] The latest Council compromise draft 2021/0106(COD) 12549/22 reads "This Regulation shall not apply to any research and development activity regarding AI systems."
[7] The latest Council compromise draft 2021/0106(COD) 12549/22 reads "Such requirements and obligations shall apply irrespective of whether the general purpose AI system is placed on the market or put into service as a pre-trained model and whether further fine-tuning of the model is to be performed by the user of the general purpose AI system."

presumption of causality in the AI Liability Directive[8] incentivizes AI system providers to select adequately documented AI-related code and models while protecting the autonomy of AI researchers and individual software developers. This liability proposal is adequate to address needs for general purpose AI models and components.

- "~~Such requirements and obligations shall apply irrespective of whether the general purpose AI system is placed on the market or put into service as a pre-trained model and whether further fine-tuning of the model is to be performed by the user of the general purpose AI system.~~"

Recommendation 5:
Acknowledge the unique role of open source software in AI research and development
Recent policy proposals from the European Commission recognize the unique characteristics of digital public good creation provided by the open source software community. The Cyber Resilience Act exempts non-commercial open source software development from scope[9], as does the Product Liability Directive proposal[10].

Similar clarity in the AI Act would provide needed certainty for open source communities that rely on volunteer contributions to build a global commons. By tying this certainty to open source AI when it is not put into service as such, certainty can be achieved in a way that does not provide perverse incentive to open source AI systems that otherwise would be subject to the

[8] COM(2022) 496: Recital 28 "The presumption of causality could also apply to AI systems that are not high-risk AI systems because there could be excessive difficulties of proof for the claimant. For example, such difficulties could be assessed in light of the characteristics of certain AI systems, such as autonomy and opacity, which render the explanation of the inner functioning of the AI system very difficult in practice, negatively affecting the ability of the claimant to prove the causal link between the fault of the defendant and the AI output. A national court should apply the presumption where the claimant is in an excessively difficult position to prove causation, since it is required to explain how the AI system was led by the human act or omission that constitutes fault to produce the output or the failure to produce an output which gave rise to the damage. However, the claimant should neither be required to explain the characteristics of the AI system concerned nor how these characteristics make it harder to establish the causal link."

[9] COM(2022) 454: Recital 10 "In order not to hamper innovation or research, free and open-source software developed or supplied outside the course of a commercial activity should not be covered by this Regulation. This is in particular the case for software, including its source code and modified versions, that is openly shared and freely accessible, usable, modifiable and redistributable. In the context of software, a commercial activity might be characterized not only by charging a price for a product, but also by charging a price for technical support services, by providing a software platform through which the manufacturer monetises other services, or by the use of personal data for reasons other than exclusively for improving the security, compatibility or interoperability of the software."

[10] COM(2022) 495: Recital 13 "In order not to hamper innovation or research, this Directive should not apply to free and open-source software developed or supplied outside the course of a commercial activity. This is in particular the case for software, including its source code and modified versions, that is openly shared and freely accessible, usable, modifiable and redistributable. However where software is supplied in exchange for a price or personal data is used other than exclusively for improving the security, compatibility or interoperability of the software, and is therefore supplied in the course of a commercial activity, the Directive should apply."

regulation. This clarity could be provided by adding a similar Recital and by striking reference in Article 3(1b)[11] to open source software.

Along these lines, some of the text presented in the Parliament Committee on Legal Affairs report dated September 12, 2022[12] is encouraging. Below that text is included, with ***additions made in italicized bold*****,** ~~deletions crossed out~~, and *notes in italics.*

- Recital 57b: "Open Source software licences allow users to run, copy, distribute, study, change and improve software freely. By default the use of Open Source software in this manner attributes liability to the user, whereas when a provider provides Open Source software commercially under a Software as a Service (SaaS) or Professional Services model, then the provider may retain the liability instead of the user. Research by the European Commission shows that Open Source software contributes between €65bn - €95bn to the European Union's GDP, and provides significant growth opportunities for the Union economy. Open Source providers should be able to adopt the same economic model for AI systems. Hence, the provisions of this Regulation should not apply to Open Source AI systems until those systems are put into service. To ensure that AI systems cannot be put into service without complying with this Regulation, when an Open Source AI System is put into service, the obligations associated with providers should be transferred to the person putting the system into service."
- Article 2(4a): "This regulation shall not apply to Open Source AI systems until those systems are put into service or made available on the market ***as a high risk AI system or*** in return for payment, regardless of if that payment is for the AI system itself, the provision of the AI system as a service, or the provision of technical support for the AI system as a service." *Additions add clarity that putting OSS AI systems into service as part of a high risk system is within scope regardless of whether or not it is a commercial activity, and draw from text considered by Axel Voss: "This regulation shall not apply to Free or Open Source AI software until it is put into service or made available as part of a high-risk AI system, as part of the provision of a high-risk AI system as a service, or the provision of technical support for a high-risk AI system as a service."*
- Article 3(1b): "'open source AI systems' means AI systems, including test and training data, or trained models, distributed under open licenses."
- Article 23a Conditions for other persons to be subject to the obligations of a provider (1d): "they place on the market or make available on the market, with or without modification***, a high-risk AI system derived from Free and Open Source AI software, or otherwise*** ~~and~~ in return for payment**,** an Open Source AI system, an AI system derived from an Open Source AI system, or Technical Support Services for any such Open Source AI systems;" *With essential clarity provided by text considered by Axel*

[11] The latest Council compromise draft 2021/0106(COD) 12549/22 reads "'general purpose AI system' means an AI system that - irrespective of how it is placed on the market or put into service, including as open source software - is intended by the provider to perform generally applicable functions such as image and speech recognition, audio and video generation, pattern detection, question answering, translation and others; a general purpose AI system may be used in a plurality of contexts and be integrated in a plurality of other AI systems;"

[12] JURI Committee report, September 12, 2022, 2021/0106(COD)

Voss, adding (2): "***Where the circumstances referred to in paragraphs 1(d) occur, developers of the Free and Open Source AI software shall not have any duty to collaborate towards the provider***."

. . .

We are optimistic that with these five revisions, the AI Act can fulfill the goals of addressing risks posed to fundamental rights and safety and of promoting AI innovation in the single market. We stand ready to discuss these matters with you and answer any questions that you may have.