

Truly open AI gets a boost from industry collaboration with enterprise implications

By [George Lawton](#) December 12, 2023

SUMMARY: Efforts to build open AI are getting a boost from the AI Alliance backed by IBM, Meta, Intel, NASA, the National Science Foundation, and fifty other organizations. This could spur research into smaller, more efficient, and secure models that may be more appropriate for enterprises.



The new AI Alliance, backed by IBM, Meta, ServiceNow, NASA, dozens of universities, and over fifty founding members, promises progress on open, safe, and responsible AI. OpenAI, which many consider not that open, is not among them.



The Alliance will help to balance industry appetites for AI with a strong roster of academic and government partners from around the world with a total budget of over \$80 billion and over 100,000. Well, mostly around the world – Chinese and Russian companies and institutions were notably absent from the map of partner locations.

Top goals of the collaboration include:

- Developing benchmarks and evaluation standards, tools, and other resources that enable the responsible development and use of AI systems at global scale, including creating a catalog of vetted safety, security and trust tools.
- Responsibly advance the ecosystem of open foundation models with diverse modalities, including highly capable multilingual, multi-modal, and science models that can help address society-wide challenges in climate, education, and beyond.
- Foster a vibrant AI hardware accelerator ecosystem by boosting contributions and adoption of essential enabling software technology.
- Support global AI skills-building and exploratory research. Engage the academic community to support researchers and students in learning and contributing to essential AI model and tool research projects.
- Develop educational content and resources to inform public discourse and policymakers on AI benefits, risks, solutions, and precision regulation.
- Launch initiatives that encourage open development of AI in safe and beneficial ways, and host events to explore AI use cases and showcase how Alliance members are using open technology in AI responsibly and for good.

Open progress

Some industry leaders expect that the AI Alliance will foster the kind of cooperation and growth made possible by the widespread adoption of open-source practices. Steven Huels, General Manager Artificial Intelligence Business & Head of Artificial Intelligence Product and Strategy at Red Hat, says:

In the years prior to open source adoption, companies innovated in silos and technologies that lacked transparency and interoperability. That's roughly where we are at now with AI models and tools. The Alliance will create a catalog of vetted safety, security and trust tools. Additional focus areas for the Alliance are to enable an ecosystem of open foundation models, ensure a diverse ecosystem of AI hardware, and help to close the AI skill gap ensuring that the academic community benefits from the Alliance's thought leadership directly.

We need to identify areas of common need (both software and hardware) where organizations can engage in order to make progress on those shared problem spaces driving transparency and interoperability where possible. As in the early days of any technological paradigm shift, there is ambiguity—one of our primary goals for our members is to collaborate to develop educational content for policymakers on benefits, risks, solutions, and precision regulation for AI. This will help to flush ambiguity out of the global system so that our industry more clearly understands the landscape and is able to safely innovate at their own speed.

Even if many of the foundation models are not formally released as open-source, many of the supporting tools are already open-source. For example, several open-source projects under the Linux Foundation AI and data group already focus on explainability, privacy, adversarial robustness (security), fairness, and evaluation. Red Hat is also active in this space, founding the TrustyAI project and integrating it with their OpenShift AI hybrid cloud MLOps platform.

Open questions

The alliance does raise some interesting questions. Amanda Brock, CEO of OpenUK, a non-profit open collaboration advocate in the UK, observes that open source has played a seminal role in the evolution of software in general. Today, nearly 96% of all software has open source software dependencies. However, these voices and experiences have been largely ignored in the consultation about AI to date.

There are multiple issues at stake here. The EU may potentially carve open source out of certain requirements of AI legislation. It is not entirely clear what 'open source' will mean in this context. Brock believes that the Alliance can be critical in building this understanding if it is expanded to offer membership beyond the initial fifty.

Of particular note is that the existing Open Source Definition (OSD) as it applies to software means there is the ability for anyone to use the code for any purpose. However, licenses like the Meta's Llama Community License include an acceptable use policy which obviates this definition of OSD in the

Premier partner newsfeed

- Salesforce's Recommendations for the Belgian Presidency of the Council of the EU
- Life sciences leaders invest in process mining, transform business intelligence into actionable insight, Everest says
- Integrated vs Dedicated GPUs
- 5 ways companies can promote more sustainable AI
- Drive Operational Excellence in Presales with IFS Cloud
- Introducing the Sticky Notes extension by Zoho Notebook
- The ASUG Talks Podcast: Tech Connect Special Edition: Learning SAP Closing the Skills Gap, and Facing the Future
- Workday Named a Leader in the 2023 Gartner® Magic Quadrant™ for Financial Planning Software for Second Year in a Row
- Experience in the age of AI
- How to Calculate Your Marketing Funnel Metrics & Budget
- Confluent placed as a 'Niche Player' in the 2023 Gartner® Magic Quadrant™ for Data Integration Tools
- Women in Tech Luncheon: Expanding Your Impact—Showcasing Mindset, Tenacity, and Action
- Back to the Future: Should Service Execution begin with the handover from Engineering?
- Slaying silos and connecting the dots for services delivery success
- Why your growing company may need a Chief of Staff (and how to find the right one)
- Stéphane Nicoll on Contributing to and Growing Communities
- Six Common Reasons to Change Fund Accounting Software

[Subscribe To Our FREE Newsletter](#)

Samsara's data platform even more acute with EV transition

Latest Conversations

- Salve Regina Samarita Insigne: Building the best teams involves fostering a culture of collaboration, communication, and mutual respect. Clearly define roles and expectations, ensuring each team member understands their...
[How to build the best teams? Slack research highlights the top five work personas](#) - 8 hours ago
- glawton: The last link was a bit fickle, here is a better one <https://www.kvinsfuneralhome...>
[Final thoughts - in memory of Simon Mark Hughes, hallucination research pioneer](#) - 1 week ago
- Jack Vaughan: Well-considered analysis, Neil – and Melanie. Thanks! My 2 cents: With their admitted faults, LLMs are a step up from Transformers, which were a step up (again, with admitted faults) from clunky...
[Large Language Models versus human intelligence - where do we stand now?](#) - 1 week ago
- cliveb: The Big Four+ are sowing together the cloud apps and products, competing with the infosys and Wipro of the world. IMO because ERP vendors have removed buyer customization for SaaS economics based...
[Enterprise hits and misses - Sam Altman punted from OpenAI but finds happy landings at Microsoft. Enterprises react, while retailers ramp up holiday hopes](#) - 2 weeks ago
- Jon Reed: Nice segue, where we agree is that Bret Taylor is interesting on several levels, as a savvy enterprise technologist, responsible AI in a more day-to-day way that OpenAI needs, and - as you say - ...
[Enterprise hits and misses - Sam Altman punted from OpenAI but finds happy landings at Microsoft. Enterprises react, while retailers ramp up holiday hopes](#) - 2 weeks ago

Go to topics

diginomica

traditional sense of the term. Instead, it is commercially restricted, requiring a commercial license to be obtained at 700 million users. Meta has not elaborated on those commercial terms, but their very existence stops it from being open source in the true sense of the term.

Additionally, AI introduces new considerations aside from the code itself that include properties that are unique to AI. Brock explains:

“ We are looking not to the software of the past, but to the AI of the future, and in this context, we must consider applying openness to weights, models, and data - as well as to the code. That means that there needs to be a better understanding of the different levels of openness and components that can be opened up and to the place each has in AI. Of course, the level of risk is different in each, and therefore, the regulatory requirements differ.

Many gaps

It's also important to note that there are many gaps in the Alliance's focus area, such as in data provenance and the ethical practices used for training models. George Davis, CEO of Frame AI, notes:

“ It's a fair callout that those forms of openness are not directly included in AI Alliance's focus areas. The supply chain that creates AI is more complex and labor intensive than the software industry has dealt with previously, and I know many researchers (including myself as well as signatories at AI Alliance organizations) would love to see more leadership and targeted advocacy for both creators' rights and labor rights within that supply chain.

However, despite these gaps, Davis believes that open research creates more room to ask these types of questions even without specific advocacy. Additionally, watchdogs can inspect and run open models independently to find evidence relating to source data. And open research projects like Databricks' April release of Dolly have provided models of how ethically sourced data can be used efficiently to advance the state of the art.

Improving enterprise offerings

Additionally, Davis expects the alliance to promote innovation through specialized models that are better suited for many practical use cases:

“ A lot of public attention has been focused on the progress of cutting edge, expensive, generalist foundation models, such as GPT-4 and Gemini. But much of the economic value AI will generate must come through specialized, less expensive models that can be deployed broadly without rerouting the power grid. The open research community is well-suited and, so far, disproportionately responsible for the exploration of these specialized models. Promoting open research will lead to a much more robust AI ecosystem where users can pick the right tool for a task, rather than paying top dollar for a model that can plan a menu or write a sonnet with equal facility.

Davis also believes the alliance will publicize best practices for building secure AI and the tools to make doing so more attainable for more participants. For example, Meta followed up on the launch of the Alliance with the release of Purple Llama, a suite of practices and tools for building safer AI models.

Davis predicts:

“ I think in the coming months, we'll see an increase in projects like this from the open community, many of which may plug holes in the threat models being considered by closed labs today. I believe the open community will demonstrate to regulators that we can have a more robust security/safety ecosystem when the rules permit a wide variety of actors to participate.

My take

It is worth noting that the AI Alliance work is not technically advocating for open-source, data transparency, or development process transparency, which sometimes gets conflated with other kinds of open things. They are, in fact, advocating for open and transparent innovation, open development, open science, open technologies, open community, open foundation models, and many other open things. The IBM press release includes the word 'open' 85 times.

Some of the open-source contributions do not necessarily come with transparency about what data they were trained on or the labor practices involved. Transparent was only mentioned once in the actual press release and then another seven times in the follow-up quotes.

The transformation of open source into other kinds of openness for promoting innovation is not necessarily a bad thing.

As Davis notes:

“ In my view, there are many forms of openness, but they lead into each other even when pursued independently.

Regulators worldwide are all grappling with balancing the obvious benefits that have accrued from traditional open-source models against the tradeoffs of enabling bad actors and geopolitical adversaries. Including lots of academic institutions that may guide the development of responsible and ethical AI will make a lot of sense. Only time will tell how this all pans out.

Brock believes we need to think about the nuances of shades of openness:

“ The most efficient regulatory response will be a joined-up international one, and that joined-up response needs to be built on mature and cohesive input from industry, academia, and the representative bodies of open technology - not just open source software. Understanding the different opens (software, hardware, data, of old and now the weights and models of AI) and that what is shared as open may include 'shades of openness' will be crucial.

We need to ensure that regulation is appropriate to the levels of openness applied. For example, a truly open source LLM like Falcon, with its software on an Open Source Initiative approved license, may well be in a different regulatory position from one like Llama distributed on the Llama Community License with commercial restrictions, which have very different implications. If we don't build this understanding into regulation from the beginning, then it will be fundamentally

flawed, and we'll be doing the equivalent of regulating a bicycle and a car with the same regulations as they are both vehicles.

Image credit - Pixabay

Get your weekly enterprise AI digest

Complete the form below to receive the top enterprise AI stories from diginomica, every week.

First Name*

Last Name*

Email*

Read more on: [Machine intelligence and AI](#) | [Ethics](#) | [Regulation](#) | [Open source](#) | [Audio](#)

diginomica and the diginomica logo are trademarks of diginomica Limited.
© DIGINOMICA LIMITED AND ITS LICENSORS 2013- 2023

[Cookie settings](#)
Developed by [BRAINSUM](#)

[Subscribe To Our FREE Newsletter +](#)