

D E E P D I V E : A I



2023 REPORT UPDATE



open source
initiative®

Introduction

Stefano Maffulli



As the leading voice on the policies and principles of Open Source, the Open Source Initiative (OSI) has played a pivotal role in shaping the landscape of the modern software ecosystem through its stewardship of the Open Source Definition. This definition, which has been the cornerstone for over 25 years, has facilitated seamless collaboration, information sharing, research endeavors, and commercial activities within the realm of software development. It has become the bedrock upon which entire generations of technology and regulatory frameworks have been built, fostering a common understanding that transcends boundaries.

Building on this rich legacy, the Open Source Initiative is now bringing this same clarity to the next generation of technological systems: Artificial Intelligence (AI). Recognizing the transformative impact of AI on society, the OSI is actively engaged in a multi-stakeholder Deep Dive: AI initiative to define what “Open Source AI” means. As part of this open process, the OSI has brought together a diverse cohort of leaders – from law, academia, non-governmental organizations (NGOs), enterprise, and the open source community – to be part of the Deep Dive: AI webinar series.

The webinar series contemplated 17 sessions featuring insights from 34 speakers from around the world. They have served as a platform for thought leaders to delve into the pressing issues surrounding AI, offering a spectrum of opinions on the challenges faced by society in the realm of Artificial Intelligence. The diversity of perspectives presented in these sessions not only underscores the complexity of the issues at hand but also hints at potential solutions that can guide the responsible development and utilization of AI technologies.

The Open Source Initiative's foray into the realm of AI represents a natural progression in its commitment to maintaining the integrity of Open Source principles. By bringing together experts from various domains, the OSI is actively contributing to the discourse on Open Source AI, laying the groundwork for a future where the principles of openness, transparency, and collaboration continue to underpin the evolution of cutting-edge technologies for the benefit of society as a whole.

 <https://opensource.org/events/deep-dive-ai-webinar-series-2023/>

The Turing Way: Who is Building Open Source AI?

Jennifer Ding, Arielle Bennett, Anne Steele, Kirstie Whitaker, Marzieh Fadaee, Abinaya Mahendiran, David Gray Widder, Mophat Okinyi



The Turing Way proposes hosting a panel discussion centering key communities who are part of building AI today, whose contributions are often overlooked. Through a conversation with panelists from content moderation, data annotation, and data governance backgrounds, we aim to highlight different kinds of contributors whose work is critical, but are often left out of governance decisions or from benefiting from the AI value chain. We will focus on these different forms of work and how each is recognised and rewarded within the open source ecosystem, with an eye to what is happening now in the AI space. In the spirit of an AI openness that promotes expanding diverse participation, democratizing governance, and inviting more people to shape and benefit from the future of AI, we will frame a conversation that highlights current best practices as well as legal, social, and cultural barriers. We hope this multi-domain, multi-disciplinary discussion can emphasize the importance of centering the communities who are integral to AI production in conversations, considerations, and definitions of “Open Source AI.”

Webinar summary [🔗](#)

In the "Fireside Chat | Who is building open source AI?" webinar, Anne Lee Steele, the community manager of the Turing Way, facilitated a discussion among panelists Marzieh, Abinaya, and David, focusing on the often overlooked contributors to open source AI, such as content moderators, data curators, and data stewards. They emphasized the need to acknowledge and reward these contributors for their important roles and the depth of their expertise. Marzieh highlighted the significance of openness in AI research and the importance of collaboration, transparency, and accountability. Mophat advocated for legislation to protect content moderators and the recognition of content moderation as a professional job. David Widder warned about potential exploitation in the open source movement and the need to critically evaluate the role of openness in the development and implementation of AI technology. Overall, the panelists emphasized the importance of openness and cooperation in building a future with AI that benefits everyone.

Operationalizing the SAFE-D Principles for Open Source AI

Kirstie Whitaker



The SAFE-D principles (Leslie, 2019) were developed at the Alan Turing Institute, the UK's national institute for data science and artificial intelligence. They have been operationalised within the Turing's Research Ethics (TREx) institutional review process. In this panel we will advocate for the definition of Open Source AI to include reflections on each of these principles and present case studies of how AI projects are embedding these normative values in the delivery of their work.

The SAFE-D approach is anchored in the following five normative goals:

- Safety and Sustainability
- Accountability
- Fairness and Non-Discrimination
- Explainability and Transparency
- Data quality, integrity, protection and privacy

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Kirstie Whitaker, Professor David Leslie and Victoria Kwan of The Alan Turing Institute discuss the operationalization of safety principles in AI research and development within the context of open source practices. Leslie introduces the safety principles of safety, accountability, fairness, explainability, and data stewardship (SAFE-D) and explains their significance in ensuring responsible and trustworthy AI development. Kwan then demonstrates how these principles are integrated into the Turing's Research Ethics (TREx) institutional review process, emphasizing stakeholder engagement, accountability, fairness, and more. Whitaker concludes by highlighting the alignment of open source practices with the safety principles, emphasizing the importance of transparency, accountability, diversity, and data stewardship within the open source AI community, while advocating for a more inclusive, accountable, and interconnected ecosystem.

Commons-based data governance

Alek Tarkowski, Zuzanna Warso



Issues related to data governance (its openness, provenance, transparency) have traditionally been outside the scope of open source frameworks. Yet the development of machine learning models shows that concerns over data governance should be in the scope of any approach that aims to govern open-source AI in a holistic way. In this session, I would like to discuss issues such as: – the need for openly licensed / commons based data sources – the feasibility of a requirement to openly share any data used in the

training of open-source models – transparency and provenance requirements that could be part of an open-source AI framework.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Alex Tarkowski and Zuzanna Warso discuss the importance of commons-based data governance in the context of open source AI and democratizing AI development. They emphasize the need to rethink how AI datasets are managed and shared, considering the diversity of data sources and the challenges of balancing openness with privacy and fairness. The speakers propose four key principles for commons-based data set governance: sharing data as openly as possible, respecting the decisions of data subjects and creators, ensuring sustainability and fair rewards for contributors, and protecting the commons from pollution and biases. They argue that these principles can help create a roadmap for developing and maintaining open datasets in a fair and sustainable manner, ultimately benefiting the open source AI community.

Preempting the Risks of Generative AI: Responsible Best Practices for Open-Source AI Initiatives

Monica Lopez, PhD



As artificial intelligence (AI) has proliferated across many industries and use cases, changing the way we work, interact and live with one another, AI-enabled technology poses two intersecting challenges to address: the influencing of our beliefs and the engendering of new means for nefarious intent. Such challenges resulting from human psychological tendencies can inform the type of governance needed to ensure safe and reliable generative AI development, particularly in the domain of open-source content.

The formation of human beliefs from a subset of available data from the environment is critical for survival. While beliefs can change with the introduction of new data, the context in which such data emerges and the way in which such data are communicated all matter. Our live dynamic interactions with each other underpin our exchange of information and development of beliefs. Generative AI models are not live systems, and their internal architecture is incapable of understanding the environment to evaluate information. Considering this system reality with the use of AI as a tool for malicious actors to commit crimes, deception –strategies humans use to manipulate others, withhold the truth, and create false impressions for personal gain– becomes an action further amplified by impersonal, automated means.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Dr. Monica Lopez, co-founder and CEO of Cognitive Insights for Artificial Intelligence, discusses the challenges posed by the proliferation of generative AI, particularly large language models (LLMs), focusing on two key issues: the influence of AI-generated content on human beliefs and the potential for nefarious uses of the technology. Dr. Lopez emphasizes that generative AI systems lack understanding and can produce biased, fabricated, and misleading content. She proposes psychological solutions, including toning down anthropomorphic language when describing AI capabilities and integrating empirical measurements of AI’s impact on human beliefs into system audits. Additionally, she addresses the risks of misuse, emphasizing the need for transparency, fairness, accountability, and reliability in open source AI initiatives. Dr. Lopez calls for benchmarks that align with societal goals and highlights the importance of interdisciplinary collaboration in addressing these challenges while placing human control at the center of AI development and use.

Data privacy in AI

Michael Meehan



Data privacy in AI is something everyone needs to plan for. As AI technology continues to advance, it is becoming increasingly important to protect the personal information that is used to train and power these systems, and to ensure that companies are using personal information properly. First, understand that AI systems can inadvertently leak the data used to train the AI as it is producing results. This talk will give an overview of how and why this happens. Second, ensure that you have proper rights to use data fed into your AI. This is not a simple task at times, and the stakes are high. This talk will go into detail about circumstances where the initial rights were not proper, and the sometimes-catastrophic results of that. Third, consider alternatives to using real personal information to train models. One particularly appealing approach is to use the personal data to create statistically-similar synthetic data, and use that synthetic data to train your AI systems. The considerations are important to help protect personal information, or other sensitive information, from being leaked by using AI. This will help to ensure that AI technology can be used safely and responsibly, and that the benefits of AI can be enjoyed with fewer risks.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Michael Meehan emphasizes the importance of synthetic data as a means to protect personal information and sensitive data while training AI models. Using the example of a company called Everalbum, the speaker suggests that Everalbum could have avoided privacy issues if it had utilized synthetic data in its model training process. Synthetic data can legally circumvent privacy concerns because it is not derived from real, identifiable information. While acknowledging that synthetic data may sacrifice some accuracy compared to real data, Meehan underscores the crucial balance between privacy and accuracy. He stresses that these considerations extend beyond personal information and can apply to trade secrets or other sensitive data. Ultimately, the use of synthetic data can promote while minimizing risks associated with data privacy.

Perspectives on Open Source Regulation in the upcoming EU AI Act

Katharina Koerner



This presentation will delve into the legal perspectives surrounding the upcoming EU AI Act, with a specific focus on the role of open source, non-profit, and academic research and development in the AI ecosystem. The session will cover crucial topics such as defining open data and AI/ML systems, copyrightability of AI outputs, control over code and data, data privacy, and fostering fair competition while encouraging open innovation. Drawing from existing and upcoming AI regulations globally, we will present recommendations to facilitate the growth of an open ecosystem while safeguarding ethical and accountable AI practices.

Join this session for an insightful exploration of the legal landscape shaping the future of open source.

What You Will Learn in the Presentation:

- The key problems faced by open source projects under the draft EU AI Act.
- The significance of clear definitions and exemptions for open source AI components.
- The need for effective coordination and governance to support open source development.
- The challenges in implementing the R&D exception for open source AI.
- The importance of proportional requirements for “foundation models” to encourage open source innovation and competition.
- Recommendation to address the concerns of open source platform providers and ensure an open and thriving AI ecosystem under the AI Act.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Katharina Koerner, from the Tech Diplomacy Network, discusses the implications of the upcoming EU AI Act for the open source ecosystem. She highlights the significance of open source in the EU and the various initiatives supporting it, emphasizing that the EU is committed to open source as a driver of innovation and accessibility. Katharina provides an overview of the draft EU AI Act, which categorizes AI systems based on risk levels and outlined the scope of the Act, including its application to providers, deployers, importers, and distributors. She also discusses exceptions for open source, emphasizing that collaborative development and open repositories are not considered making AI systems available on the market unless they turn commercial. Furthermore, she explains the unique challenges and requirements related to foundation models within the AI Act, underscoring the need for compliance, transparency, and responsible practices in the open source AI community to prepare for the impending regulations.

Data Cooperatives and Open Source AI

Tarunima Prabhakar, Siddharth Manohar



Data Cooperatives have been proposed as a possible remediation to the current power disparity between citizens/internet users from whom data is generated and corporations that process data. But these cooperatives may also evolve to develop their own AI models based on the pooled data. The move to develop machine learning may be driven by a need to make the cooperative sustainable or to address a need of the people pooling the data. The

cooperative may consider ‘opening’ its machine learning model even if the data is not open. In this talk we will use Uli, our ongoing project to respond to gendered abuse in Indian languages, as a case study to describe the interplay between community pooled data and open source AI. Uli relies on instances of abuse annotated by activists and researchers at the receiving end of gendered abuse. This crowdsourced data has been used to train a machine learning model to detect abuse in Indian languages. While the data and the machine learning model were made open source for the beta release, in subsequent iterations the team is considering limiting the data that is opened. This is, in part, a recognition that the project is compensating for the lack of adequate attention to non anglophone languages by trust and safety teams across platforms. This talk will explore the different models for licensing data and the machine learning models built on it, that the team is considering, and the tradeoffs between economic sustainability and public good creation in each.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Tarunima Prabhakar and Siddharth Manohar from Tattle Civic Tech organization in India, discuss the challenges of reconciling open source principles with the realities of data cooperatives and machine learning in the context of addressing online abuse, particularly targeting marginalized genders in India. They explore two potential scenarios for differential access to data: one based on the purpose of use (distinguishing between academic, nonprofit, and commercial users) and the other based on the timeliness and relevance of the data. They also consider the implications of these data licensing scenarios on machine learning model licensing, highlighting the balance between openness, transparency, and preventing abuse. The webinar emphasizes the importance of recognizing the labor involved in data annotation, data justice, and the need to protect against misuse of data while enabling legitimate use cases.

Fairness & Responsibility in LLM-based Recommendation Systems: Ensuring Ethical Use of AI Technology

Rohan Singh Rajput



The advent of Large Language Models (LLMs) has opened a new chapter in recommendation systems, enhancing their efficacy and personalization. However, as these AI systems grow in complexity and influence, issues of fairness and responsibility become paramount. This session addresses these crucial aspects, providing an in-depth exploration of ethical concerns in LLM-based recommendation systems, including algorithmic bias, transparency, privacy, and accountability. We'll delve into strategies for mitigating bias, ensuring data privacy, and promoting responsible AI usage. Through case studies, we'll examine real-world implications of unfair or irresponsible AI practices, along with successful instances of ethical AI implementations. Finally, we'll discuss ongoing research and emerging trends in the field of ethical AI. Ideal for AI practitioners, data scientists, and ethicists, this session aims to equip attendees with the knowledge to implement fair and responsible practices in LLM-based recommendation systems.

Webinar summary [🔗](#)

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Rohan Rajput discusses the intersection of fairness and responsibility in Language Model (LM)-based recommendation systems. He begins by introducing LM as powerful language models trained on vast textual data, highlighting their ability to generate responses based on user prompts. Recommendation systems, such as those used by Amazon and Netflix, are then introduced as a domain within Information Retrieval Systems. The main focus of the presentation is on LM-based recommendation systems, categorized into prediction and generation tasks. Rajput emphasizes the ethical challenges these systems face, particularly in domains like education, criminology, finance, and health, due to the potential replication of biases present in training data. He discusses fairness and bias mitigation strategies, including robust data processing, algorithmic fairness, multi-objective optimization, transparency, and user control. Rajput also touches on issues like hallucination, fabricating information, and the importance of diversity and compliance. He concludes by highlighting ongoing efforts in user education, monitoring, third-party audits, community involvement, and public input, stressing that fairness in recommendation systems is an ongoing, iterative process requiring continuous improvement and vigilance.

Challenges welcoming AI in openly-developed open source projects

Thierry Carrez, Davanum Srinivas, Diane Mueller



Openly-developed open source projects are projects that are developed in a decentralized manner, fully harnessing the power of communities by going beyond open source to also require open development, open design and open community (the 4 opens). This open approach to innovation has led to creation

of very popular open source infrastructure technologies like OpenStack or Kubernetes.

With the rise of generative solutions and LLMs, we are expecting more and more code to be produced, directly or indirectly, by AI. Expected efficiencies may save millions of dollars. But at what cost? How is that going to affect the 4 opens? What are the challenges in welcoming AI in our open communities?

- This webinar will explore questions such as:
- Can AI-generated code be accepted in projects under an open source license?
- How can we expect open design processes to evolve in a AI world?
- Is it possible to avoid that the burden just shifts from code authoring to code reviewing?
- What does open community mean with AI-powered participants? Is there a risk to create a second class of community members?

Webinar summary [🔗](#)

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, a panel of open source community leaders, including Thierry Carrez, Diane Mueller, and Dims Srinivasi, explores the challenges and implications of integrating AI into open source projects. They discuss the impact of AI on the four facets of open collaboration: open source, open development, open design, and open community. The panelists express concerns about the potential influx of AI-generated contributions, putting added pressure on project reviewers and maintainers. They highlight the need for transparency, ethics, and governance in handling AI-generated code, emphasizing the importance of credit, consent, and compensation for creators. The discussion also touches on the potential benefits of AI in automating code generation, bug detection, and even collaborative design processes while raising questions about biases in AI-generated content. Ultimately, the panel envisions a future where AI plays a significant role in open source, but ethical and governance challenges need to be addressed.

Opening up ChatGPT: a case study in operationalizing openness in AI

Andreas Liesenfeld, Mark Dingemans



Openness in AI is necessarily a multidimensional and therefore graded notion. We present work on tracking openness, transparency and accountability in current instruction-tuned large language models. Our aim is to provide evidence-based judgements of openness for over ten specific features, from source code to training data to model weights and from licensing to scientific documentation and API access. The features are grouped in

three broad areas (availability, documentation, and access methods). The openness judgements can be used individually by potential users to make informed decisions for or against deployment of a particular architecture or model. They can also be used cumulatively to derive overall openness scores (tracked at <https://opening-up-chatgpt.github.io>). This approach allows us to efficiently point out questionable uses of the term “open source” (for instance, Meta’s Llama2 emerges as the least open of all ‘open’ models) and to incentivise developers to consider openness and transparency throughout the model development and deployment cycle (for instance, the BLOOMZ model stands out as a paragon of openness). While our focus is on LLM+RLHF architectures, the overall approach of decomposing openness into its most relevant constituent features is of general relevance to the question of how to define “open” in the context of AI and machine learning. As scientists working in the spirit of open research, the framework and source code underlying our openness judgements and live tracker is itself open source.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Andreas Liesenfeld and Mark Dingemans with the Center for Language Studies at Radboud University in the Netherlands discuss the “Opening Up Chat GPT” project, which aims to evaluate the openness of large language models and text generators. The project addresses the need for open source AI technology in European academia and introduces a methodology to assess the openness of these systems. The evaluation focuses on various dimensions of openness, including the availability of source code, documentation, user access, pre-training datasets, model weights, licenses, and more. Using a comparative analysis of two systems, the BLOOMZ model and Meta’s Llama2, the presentation illustrates how this evidence-based approach can differentiate between genuinely open systems and those with limited openness. The project also highlights emerging challenges in the field, such as the use of synthetic data and the need for transparency in multi-step training pipelines.

Open source AI between enablement, transparency and reproducibility

Ivo Emanuilov, Jutta Suksi



Open source AI is a misnomer. AI, notably in the form of machine learning (ML), is not programmed to perform a task but to learn a task on the basis of available data. The learned model is simply a new algorithm trained to perform a specific task, but it is not a computer program proper and does not fit squarely into the protectable subject matter scope of most open source software licences. Making available the training script or the model's 'source code' (eg, neural weights), therefore, does not

guarantee compliance with the OSI definition of open source as it stands because AI is a collection of data artefacts spread across the ML pipeline.

The ML pipeline is formed by processes and artefacts that focus on and reflect the extraction of patterns, trends and correlations from billions of data points. Unlike conventional software, where the emphasis is on the unfettered downstream availability of source code, in ML it is transparency about the mechanics of this pipeline that takes centre stage.

Transparency is instrumental for promoting use maximisation and mitigating the risk of closure as fundamental tenets of the OSS definition. Instead of focusing on single computational artefacts (eg, the training and testing data sets, or the machine learning model), a definition of open source AI should zoom in on the 'recipe', ie the process of making a reproducible model. Open source AI should be less interested in the specific implementations protected by the underlying copyright in source code and much more engaged with promoting public disclosure of details about the process of 'AI-making'.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the "Deep Dive: Defining Open Source AI" series, Ivo Emanuilov and Jutta Suksi discuss open source AI and its components with a focus on understanding the various phases and layers involved in AI development, particularly with regard to copyright issues. Their presentation highlights the hybrid nature of AI, blending data, software, and other components, leading to complex questions about intellectual property rights. It explores the idea of replicating open source principles for AI by emphasising transparency, enablement, and reproducibility as key principles. Additionally, Ivo and Jutta examine existing AI licences and their focus on data, restrictions, weights, and combinations of data, executable models, and source code. This presentation emphasises the rapidly evolving nature of AI licences and the importance of considering regulatory requirements in the development of open source AI.

Federated Learning: A Paradigm Shift for Secure and Private Data Analysis

Dimitris Stripelis



There are situations where data relevant to a machine learning problem are distributed across multiple locations that cannot share the data due to regulatory, competitiveness, security, or privacy reasons. Federated Learning (FL) is a promising approach to learning a joint machine learning model over all the available data across silos without transferring data to a centralised location. Federated Learning was originally introduced by Google in 2017 for next-word prediction on edge devices. Recently, Federated Learning has witnessed vast applicability across multiple disciplines, especially in healthcare, finance, and manufacturing.

Typically, a federated environment consists of a centralized server and a set of participating devices. Instead of sending the raw data to the central server, devices only send their local model parameters trained over their private data. This computational approach has a great impact on how traditional training of the machine and deep learning models is performed. Compared to centralized machine learning where data need to be aggregated in a centralized location, Federated Learning allows data to live at their original location, hence improving data security and reducing associated data privacy risks. When Federated Learning is used to train models across multiple edge devices, e.g., mobile phones, sensors, and the like, it is known as cross-device FL, and when applied across organizations it is known as cross-silo FL.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Dimitri Stripelis, a research scientist at the Information Sciences Institute at the University of Southern California, discusses the promising approach of federated learning for secure and private data analysis. Traditional centralized machine learning faces challenges due to data sharing restrictions imposed by various data regulations worldwide. Federated learning offers a solution where data remains decentralized, and only local model parameters are shared. This approach has applications in healthcare, data engineering, mobile/IoT devices, pharmaceuticals, finance, and more. Challenges in federated learning include computational and statistical heterogeneity, semantic heterogeneity, and data fragmentation, which impact model performance. Security and privacy concerns necessitate encryption methods like fully homomorphic encryption to protect sensitive data. Federated learning frameworks like Metis are emerging, but standardization and benchmarking are ongoing. The goal is for federated learning to become the standard for distributed AI model training across various domains.

Should OpenRAIL licenses be considered OS AI Licenses?

Daniel McDuff, Danish Contractor, Luis Villa, Jenny Lee



Advances in AI have been enabled in-part thanks to open source (OS) which has permeated ML research both in the academy and industry. However, there are growing concerns about the influence and scale of AI models (e.g., LLMs) on people and society. While openness is a core value for innovation in the field, openness is not enough and does not address the risks of harm that might exist when AI is used negligently or maliciously. A growing category of licenses are open responsible [AI licenses](#) which include behavioral-use clauses, these include high profile projects such as Llama2 (<https://ai.meta.com/llama/>) and [Bloom](#). In this proposed session the panelists would discuss whether [OpenRAIL](#) licenses should be considered as OS AI licenses.

Topics will include: Whether the definition of OS is not adequate for AI systems; Whether OS of AI systems requires open-sourcing every aspect of the model (data, model, source) and whether that is feasible; How data use requirements could be included in such a definition; and therefore, whether inclusion of behavioral use restrictions is at odds with any future definition of OS AI. In responding to these questions the panelists will discuss how the components of AI systems (e.g., data, models, source code, applications) each have different properties and whether this is part of the motivation for a new form of licensing. The speakers have their own experience of building, distributing and deploying AI systems and will provide examples of these considerations in practice.

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Daniel McDuff, Danish Contractor, Luis Villa and Jenny Lee discuss the evolution of AI licensing, particularly within the context of open source. Initially, RAIL licenses were conceived in response to concerns about AI technologies being released without proper ethical considerations, inspired by events like the Cambridge Analytica scandal. They aimed to self-regulate and establish norms within the AI community. However, as AI models became more complex, the question arose about how to protect and standardize the use of these models and data. Some argue that the term “open source” should not be diluted and that other terms like “responsible source” should be used to describe AI models with use restrictions. The tension between the desire for open access and the need for responsible use remains a key challenge in the evolving AI landscape, with calls for tooling and standardization to facilitate compliance and enforcement of use restrictions.

Copyright – Right Answer for Open Source Code, Wrong Answer for Open Source AI?

McCoy Smith



Open source has always found its legal foundation primarily in copyright. Although many codes of behavior around open source have been adopted and promulgated by various open source communities, in the end it is the license attached to any piece of open source that dictates how it may be used and what obligations a user must abide by in order to remain legally compliant.

Artificial Intelligence is raising, and will continue to raise, profound questions about how copyright law applies — or does not apply — to the process of ingesting training content, processing that content to extract information used to generate output, what that information is, and the nature of the output produced.

Much debate, and quite a bit of litigation, has recently been generated around questions raised by the input phase of training Artificial Intelligence, and to what extent the creators of materials used in that input phase have any right — morally or legally — to object to that training. At the same time, whether or not the output of AI can be the subject matter of copyright, or patent, protection is also being tested in various jurisdictions — with clashing results. What occurs between input and output remains an unresolved issue — and whether there is any legal regime that can be used to guarantee that legal, normative rules can control how those processes are used exist in the way that copyright, and copyright licensing, do so in open source at present.

The presentation will discuss these issues in depth with a lens toward testing whether copyright — or any other intellectual property regime — really can be useful in keeping AI “open.”

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, McCoy Smith provides an overview of open source licensing, copyright, and its application to artificial intelligence (AI) systems. McCoy discusses the history of open source licensing and how it has primarily relied on copyright law. He explores the challenges and debates surrounding AI systems and copyright, especially concerning training data, weights, and vectors generated by these systems. McCoy speculates on potential future developments, including legislative solutions, contract-based approaches, and public pledges as means to address the evolving landscape of open source AI. Finally, he highlights the complexities and legal ambiguities arising from the interplay of copyright law and AI technology.

Should we use open source licenses for ML/AI models?

Mary Hardy



Open source AI models are exponentially increasing in number and the variety of open source licenses chosen is substantial. Can all OSI-approved licenses be used uniformly to fit the various components of AI?

During the session, open source attorney Mary Hardy will explore questions present and future about open ML model licenses, including:

Why is AFL-3.0 so popular?

What about Apache-2.0? GPL-2.0/3.0?

What are the implications of licensing modifications under a different OS license than the checkpoint used as a basis?

Is a new license that explicitly considers ML model weights needed?

Webinar summary

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Mary Hardy, corporate counsel at Microsoft, discusses the challenges and complexities surrounding the application of open source licenses to machine learning models and components. She highlights the ambiguities and contradictions in existing licenses like GPL 3.0, GPL 2.0, and the Academic Free License 3.0 (AFL) when it comes to defining what constitutes the “preferred form for modification” and how these licenses may not align with the evolving nature of machine learning models and datasets. Hardy also mentions instances where licenses are specifically tailored to include machine learning components, but these adaptations can lead to compatibility issues and restrictions on the use of trained models and weights. Hardy questions whether creating new, highly specific licenses is a solution or if it would exacerbate the problem of license proliferation and incompatibility in the open source AI community.

Covering your bases with IP indemnity

Justin Dorfman, Tammy Zhu, Samantha Mandell



When working with LLM providers that don't have their models public (Anthropic, OpenAI, etc.), it's near impossible to know if any Copyleft code has been trained upon. So how do you bring AI developer tools to the market without risking legal jeopardy?

I asked Sourcegraph's head of Legal, Tammy Zhu, to teach me how we protect ourselves from failing to comply with attribution requirements.

Webinar summary [🔗](#)

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, three speakers from Sourcegraph address key topics related to AI coding assistants and intellectual property (IP) indemnity. Justin Dorfman, the Open Source Community Manager at Sourcegraph, introduces the company's AI code assistant, Cody, emphasizing its role in helping developers become more productive and write better code. Tammy Zhu, VP of Legal, discusses the legal risks associated with AI coding tools, primarily focusing on copyright risks and explaining how code snippets generated by AI tools are unlikely to be copyrightable. She also provides insights into the current litigation landscape and suggests strategies to mitigate risks. Samantha Mandel, Commercial Counsel, delves into the importance of IP indemnity in AI tools, highlighting that indemnity should benefit the end user. She outlines key elements of a comprehensive IP indemnity clause and compares Sourcegraph's approach to competitors, emphasizing Sourcegraph's commitment to protecting customers from potential legal concerns. Overall, Sourcegraph emphasizes its dedication to open source, AI, and user protection while encouraging the evolution of open source AI.

The Ideology of FOSS and AI: What “Open” means relating to platforms and black box systems

Mike Nolan



The initial conception of Free and Open Source Software was developed during a time where software was bundled into discrete packages to be run on machines owned and operated by a single individual. The initial FOSS movement utilized licensing and copyright law to provide better autonomy and control over these software systems. Now, our software systems often operate as platforms, monopolizing access between networks and resources and profiting greatly through that monopoly.

In this talk, listeners will learn more about the ideological foundations of FOSS and the blindspots that have developed in our community as software has transitioned from individual discrete packages into deeply interconnected systems that gate access to critical resources for many. We will delve into what autonomy might mean in a world where the deployment of technology inherently affects so many. Finally, we will observe the flaws in conventional open source approaches to providing autonomy and what other tools we may have at our disposal to ensure better community governance of this increasingly pervasive technology.

Webinar summary [🔗](#)

In this webinar hosted by the Open Source Initiative as a part of the “Deep Dive: Defining Open Source AI” series, Mike Nolan discusses the history and evolution of free and open-source software, examining how the concept applies to emerging technologies like AI. Mike references a recent influential paper titled “Open for Business: Big Tech Concentrated Power in the Political Economy of Open AI.” The presentation is divided into four key sections: a historical overview of open source, the changing nature of software from discrete applications to complex platforms, the components and spectrum of openness in AI, and an analysis of the pros and cons of open AI. It also touches upon the claims made regarding open AI, its impact on innovation, and the changing landscape of software development. Ultimately, Mike emphasizes that open source does not guarantee democratization or meaningful competition in AI, highlighting the need for a critical evaluation of its value in the context of AI deployment and regulation.

Acknowledgements

The Deep Dive: AI webinar series was made possible thanks to the authors, sponsors, and OSI's individual members.

Authors

Abinaya Mahendiran, Alek Tarkowski, Andreas Liesenfeld, Anne Steele, Arielle Bennett, Daniel McDuff, Danish Contractor, Davanum Srinivas, David Gray Widder, Diane Mueller, Dimitris Stripelis, Ivo Emanuilov, Jennifer Ding, Jenny Lee, Justin Dorfman, Jutta Suksi, Katharina Koerner, Kirstie Whitaker, Luis Villa, Mark Dingemans, Mary Hardy, Marzieh Fadaee, McCoy Smith, Michael Meehan, Mike Nolan, Monica Lopez, Mophat Okinyi, Rohan Singh Rajput, Samantha Mandell, Siddharth Manohar, Tammy Zhu, Tarunima Prabhakar, Thierry Carrez, Zuzanna Warso.

Organizer



**open source
initiative**[®]

Supporter



GitHub



Google Open Source

Platinum Sponsor



Gold Sponsor



Bronze Sponsor

