OPEN (FOR BUSINESS): BIG TECH, CONCENTRATED POWER, AND THE POLITICAL ECONOMY OF OPEN AI

David Gray Widder
Carnegie Mellon University
Pittsburgh, United States.

Meredith Whittaker
Signal Foundation
San Francisco, United States.

Sarah Myers West
AI Now Institute
New York, United States.

August 16, 2023

ABSTRACT

This paper examines 'open' AI in the context of recent attention to open and open source AI systems. We find that the terms 'open' and 'open source' are used in confusing and diverse ways, often constituting more aspiration or marketing than technical descriptor, and frequently blending concepts from both open source software and open science. This complicates an already complex landscape, in which there is currently no agreed on definition of 'open' in the context of AI, and as such the term is being applied to widely divergent offerings with little reference to a stable descriptor.

So, what exactly is 'open' about 'open' AI, and what does 'open' AI enable? To better answer these questions we begin this paper by looking at the various resources required to create and deploy AI systems, alongside the components that comprise these systems. We do this with an eye to which of these can, or cannot, be made open to scrutiny, reuse, and extension. What does 'open' mean in practice, and what are its limits in the context of AI? We find that while a handful of maximally open AI systems exist, which offer intentional and extensive transparency, reusability, and extensibility– the resources needed to build AI from scratch, and to deploy large AI systems at scale, remain 'closed'—available only to those with significant (almost always corporate) resources. From here, we zoom out and examine the history of open source, its cleave from free software in the mid 1990s, and the contested processes by which open source has been incorporated into, and instrumented by, large tech corporations. As a current day example of the overbroad and ill-defined use of the term by tech companies, we look at 'open' in the context of OpenAI the company. We trace its moves from a humanity-focused nonprofit to a for-profit partnered with Microsoft, and its shifting position on 'open' AI. Finally, we examine the current discourse around 'open' AI–looking at how the term and the (mis)understandings about what 'open' enables are being deployed to shape the public's and policymakers' understanding about AI, its capabilities, and the power of the AI industry. In particular, we examine the arguments being made for and against 'open' and open source AI, who's making them, and how they are being deployed in the debate over AI regulation.

Taken together, we find that 'open' AI can, in its more maximal instantiations, provide transparency, reusability, and extensibility that can enable third parties to deploy and build on top of powerful off-the-shelf AI models. These maximalist forms of 'open' AI can also allow some forms of auditing and oversight. But even the most open of 'open' AI systems do not, on their own, ensure democratic access to or meaningful competition in AI, nor does openness alone solve the problem of oversight and scrutiny. While we recognize that there is a vibrant community of earnest contributors building and contributing to 'open' AI efforts in the name of expanding access and insight, we also find that marketing around openness and investment in (somewhat) open AI systems is being leveraged by powerful companies to bolster their positions in the face of growing interest in AI regulation. And that some companies have moved to embrace 'open' AI as a mechanism to entrench dominance, using the rhetoric of 'open' AI to expand market power while investing in 'open' AI efforts in ways that allow them to set standards of development while benefiting from the free labor of open source contributors.

Keywords artificial intelligence · open source · political economy · policy · AI · Big Tech · privacy · data · competition

1 Introduction

This paper examines ‘open’ AI in the context of recent attention to open and open source AI systems. We find that the terms ‘open’ and ‘open source’ are used in confusing and diverse ways, often constituting more aspiration or marketing than technical descriptor, and frequently blending concepts from both open source software and open science. This complicates an already complex landscape, in which there is currently no agreed on definition of ‘open’ in the context of AI, and as such the term is being applied to widely divergent offerings with little reference to a stable descriptor.[1]

As a rule, ‘open’ refers to systems that offer transparency, reusability, and extensibility–they can be scrutinized, reused, and built on. But what this means in practice varies widely.[2] AI systems are composed of multiple interconnected components. As Irene Solaiman notes, AI labeled ‘open’ exists on a long gradient, with the term used to describe both systems that offer minimal transparency and reusability attributes–say, making only one of its many components open for scrutiny or reuse–alongside those that offer maximal transparency, reusability, and extensibility.[3]

AI systems don’t operate like traditional software – they require distinct development processes and rely on specialized and costly resources currently pooled in the hands of a few large tech companies. Even so, many of the original promises of the open source movement, which were made in reference to open source software, are now being projected onto ‘open’ AI. From the promise that open source could democratize software development, to the view that many eyes on open code could ensure it did what it said it did and was free of vulnerabilities, to the perspective that open source levels the playing field, allowing the most innovative to triumph. Open source software did many of these things, to varying degrees. But ‘open’ AI is a different story.

Some ‘open’ AI does extend the small circle of those who have access to well documented, already built large AI systems, expanding it from a handful of large companies to a broader developer and academic ecosystem, enabling them to examine, reuse, and build on top of these systems. Some also enable beneficial auditing and oversight that would not be possible otherwise. But even maximalist varieties of ‘open’ AI don’t democratize or extend access to the resources needed to build AI from scratch–during which highly significant ‘editorial’ decisions are made. This is a key distinction between ‘openness’ in AI and in other kinds of software. Nor does openness lower the cost of deploying large AI systems at scale, thus limiting the practical definition of ‘reuse.’ Finally, even maximal transparency in the context of AI systems doesn’t equal the same level of scrutability that open code and clear documentation do with traditional software. In the case of AI, code and documentation can’t tell you exactly how a model will perform in a given context, or enable you to predict the system’s emergent properties.[4]

So, what exactly is ‘open’ about ‘open’ AI, and what does ‘open’ AI enable? To better answer these questions we begin this paper by looking at the various resources required to create and deploy AI systems, alongside the components that comprise these systems. We do this with an eye to which of these can, or cannot, be made open to scrutiny, reuse, and extension. What does ‘open’ mean in practice, and what are its limits in the context of AI?

We find that while a handful of maximally open AI systems exist, which offer intentional and extensive transparency, reusability, and extensibility– the resources needed to build AI from scratch, and to deploy large AI systems at scale, remain ‘closed’—available only to those with significant (almost always corporate) resources. From here, we zoom out and examine the history of open source, its cleave from free software in the mid 1990s, and the contested processes by which open source has been incorporated into, and instrumented by, large tech corporations. As a current day example of the overbroad and ill-defined use of the term by tech companies, we look at ‘open’ in the context of OpenAI the company. We trace its moves from a humanity-focused nonprofit to a for-profit partnered with Microsoft, and its shifting position on ‘open’ AI. Finally, we examine the current discourse around ‘open’ AI–looking at how the term and the (mis)understandings about what ‘open’ enables are being deployed to shape the public’s and policymakers’ understanding about AI, its capabilities, and the power of the AI industry. In particular, we examine the arguments being made for and against ‘open’ and open source AI, who’s making them, and how they are being deployed in the debate over AI regulation.

Taken together, we find that ‘open’ AI can, in its more maximal instantiations, provide transparency, reusability, and extensibility that can enable third parties to deploy and build on top of powerful off-the-shelf AI models, and make some forms of auditing and oversight possible. But even the most open of ‘open’ AI systems do not, on their own, ensure democratic access to or meaningful competition in AI, nor does openness alone solve the problem of oversight

[1] The Open Source Initiative has recently offered a definition, which is not yet widely adopted: Maffulli.

[2] Liesenfeld, Lopez, and Dingemanse, “Opening up ChatGPT."

[3] Goeva, Stoudt, and Trisovic, “Toward Reproducible and Extensible Research."

[4] Lipton, “The Mythos of Model Interpretability."; Poursabzi-Sangdeh et al., “Manipulating and Measuring Model Interpretability."; Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead."

and scrutiny. While we find a vibrant community of earnest contributors building and contributing to ‘open’ AI efforts in the name of expanding access and insight, we also find that marketing around openness and investment in (somewhat) open AI systems is being leveraged by powerful companies to bolster their positions in the face of growing interest in AI regulation. And that some companies have moved to embrace ‘open’ AI as a mechanism to entrench dominance, using the rhetoric of ‘open’ AI to expand market power, and investing in ‘open’ AI efforts in ways that allow them to set standards of development while benefiting from the free labor of open source contributors.

2 Defining ‘open’ AI

Let’s start with definitions. Openness in AI is an encumbered concept. This is in part because both the definition of AI itself, and of what ‘open’ means in the context of emergent, complex systems like AI, are contested and not clearly defined.

We can start with the instability of the term AI, which across its over 70 year history has been applied to a grab bag of approaches not as a technical term of art, but more as marketing and a signifier of aspirations.[5] Currently, the term is most often applied to describe large, resource intensive machine learning systems, with generative AI currently wearing the crown as the trendiest subvariant. However, given the flexibility of the term AI, it is also applied as marketing to a deluge of computational products that as a rule, due to their proprietary nature, cannot be examined or validated.

Because large and generative AI systems are those that most clearly perturb the boundaries and question the traditional ideologies of open source and open science, we focus primarily on those systems in this paper. When we use the term AI, we are using it to refer to these large systems.

The lack of an agreed definition of AI, and the concomitant difficulty in defining what counts as ‘part of’ a multi-part AI system makes deciding what ’open’ AI is very difficult. Indeed, there is currently no agreed on definition of ‘open’ or ‘open source’ AI, even as attention to the topic has exploded. A nascent process led by the Open Source Initiative (OSI) is beginning to discuss a what a definition might look like, exploring contentious issues like whether potentially sensitive training data should be made open, whether offering powerful models for reuse is safe, and whether restricting the use of ‘open’ models to ‘safe’ or ‘ethical’ domains is acceptable under the banner of ‘open source,’ or not.[6] What is clear to the OSI, and to other observers and these authors, is that the traditional definition of Open Source, written to apply to software at a very different time in tech industry history, does not cover AI.

Even so, the ideology and assumptions about open source that were forged decades ago in the context of traditional software are being mapped on to ’open’ AI, animating many misassumptions and hyperbolic claims about what it does, and does not, enable.

This is the landscape within which we’re writing, and it has shaped our choice of terms.

Unless we’re discussing a specific license or quoting a claim made verbatim, we avoid the term ‘open source.’ Instead, we use the term ’open’ AI to connote AI systems in which one or many components are offered transparently, reusably, or in ways that allow third parties to extend and build on top. In practice, this mixes definitions from traditional open source and open science–and in doing so, we are reflecting the currently muddy application of openness to AI. In making this choice we are not working to apply the spirit of the traditional definition of open source to AI efforts. We are explicitly including contested, shallow, and borderline dishonest claims of ’openness’.[7] Like the claims made by Meta about its recently released LLaMA-2 model, which has been labeled open source by the company, even as the Open Source Initiative has stated that it is not open source.[8] In choosing to include specious and shallow appellations of ‘open’ under our term, we are in no way endorsing openwashing, nor moving to take Meta (and others) at face value. Instead, because we are focused on the role that the rhetoric of openness is playing in the public’s and policy maker’s understanding of AI and the regulatory debates surrounding it, we recognize that both meaningful and specious claims are shaping this understanding. And that discerning the distinction between a ’legitimate’ claim to open AI, while intuitive in some ways, is ultimately difficult given the lack of fixed definitions.

Broadly, the terms ‘open’ and ‘open source’ are used in the context of AI in varying ways to refer to a range of capabilities that can be broadly bucketed as offering attributes of *transparency* - the ability to access and vet source code, documentation and data - *reusability* - the ability and licensing needed to allow third parties to reuse source code

[5] AI Now Institute, “What Is AI? Part 2, with Lucy Suchman."

[6] Maffulli, “Towards a Definition of ‘Open Artificial Intelligence.’"

[7] Liesenfeld, Lopez, and Dingemanse, “Opening up ChatGPT."

[8] Vaughan-Nichols, “Meta’s Llama 2 Is Not Open Source, " Nolan, “Llama and ChatGPT Are Not Open-Source."

and/or data - and *extensibility* - the ability to build on top of extant off-the-shelf models, 'tuning' them for one or another specific purpose. While the terms 'open' and 'open source' are used variably to refer to these attributes, in practice there are gradients of openness that offer vastly differing levels of access.[9] Some systems described as 'open' offer little more than an API or the ability to download a hosted model.[10] In these cases, many question whether the term should be applied at all, or if it is 'openwashing' systems that should be understood as 'closed'. Other more maximal versions of 'open' AI go further, offering access to the source code, underlying training data and full documentation, as well as licensing the AI system for wide reuse under terms aligned with the mandates of the Open Source Initiative's definition of 'open source'.[11]

These precise distinctions matter, in part because the ideologies of 'open' and 'open source' are currently being projected onto 'open' AI systems even when they don't fit– ideologies that were forged decades ago in the context of open source software at a time when the current tech industry was emergent, not immensely powerful. Reflecting more this vestigial ideology than the reality of 'open' AI or the AI industry, we see openness in AI increasingly referenced as something that can level the playing field and provide remedy to the concentration of power in the AI sector.[12] Similarly, in regulatory and governance conversations, the existence of 'open' or 'open source' AI systems is frequently invoked to suggest that the benefits from AI and the resources required to create such systems are (or can be made) readily available to those outside large tech companies.[13] Bruce Schneier and Jim Waldo recently made this claim explicitly, stating that, "We have entered an era of LLM democratization. By showing that smaller models can be highly effective, enabling easy experimentation, diversifying control, and providing incentives that are not profit motivated, open-source initiatives are moving us into a more dynamic and inclusive A.I. landscape."[14]

But the truth is less comforting. While maximally 'open' AI is a necessary condition of any hypothetically democratized and level AI playing field, it is not a sufficient condition.[15] When we examine the resources required to create and deploy AI from start to finish, and who actually has them, visions of democratic access and a competitive AI market fade. We find that even though there are a handful of meaningfully transparent, reusable, and extensible AI systems, these and all other 'open' AI exists within a deeply concentrated tech company landscape. With scant exceptions that prove the rule, only a few large tech corporations can create and deploy large AI systems at scale, from start to finish - a far cry from the decentralized and modifiable infrastructure that once animated the dream of the free/open source software movement.[16] Given the immense importance of scale to the current trajectory of artificial intelligence, this means 'open' AI cannot, alone, meaningfully 'democratize' AI, nor does it pose a significant challenge to the concentration of power in the tech industry.

Having a grasp on the multiplicity of definitions for 'open' and 'open source' AI, alongside what such systems do and do not enable, also matters because classifying a system as 'open' or 'open source' could have downstream effects on how it is regulated. In discussions around the EU's AI Act and other potential regulatory measures, there is active lobbying arguing for broad exemptions for open source AI.[17] Such claims are premised, in part, on assertions that 'open' AI enables more democratic, innovative and competitive access to large and powerful AI systems that regulation would presumably inhibit. Notably, these assertions are often made aggressively by Big Tech firms.

Observing that the most strident voices lobbying for open source exemption from regulatory scrutiny came from many of the same tech companies that would, by the logic of these arguments, be challenged by such 'open' AI, is what motivated us to write this paper. Broad claims that 'open' AI enables more democratic access to large and powerful AI systems are, we argue, severely limited. In fact, evidence suggests that without significantly checking the power of large tech companies, 'open' AI and a regulatory environment shaped by the views of the companies lobbying for open source exemptions, could serve to further concentrate, rather than challenge, Big Tech power.

Note: we do **not** take a position here on whether openness in AI is 'good' or 'bad' more generally. We recognize that authentically transparent, reusable, and extensible 'open' AI can provide valuable insights and make possible various accountability practices, while enabling those with the resources to deploy and build on top of pre-built AI components.

[9] Solaiman, "The Gradient of Generative AI Release."

[10] Liesenfeld, Lopez, and Dingemanse, "Opening up ChatGPT."

[11] The Open Source Initiative maintains a definition for open source that requires the following criteria be filled: free redistribution, source code, provisions allowing modifications and derived work, several non-discrimination and non-restriction provisions, and technological neutrality, among other elements. https://opensource.org/osd/

[12] Mozilla, "Introducing Mozilla.Ai."; Zhang and Rai, "AI Presents Nearly Level Playing Field for Startups and Big Tech."

[13] Clegg, "Openness on AI is the Way Forward."

[14] Schneier and Waldo, "Big Tech Isn't Prepared for A.I.'s Next Chapter."

[15] More than one of the authors of this paper questions whether AI can ever be democratized.

[16] Thanks to Chris Kelty for this point.

[17] Karabus, "GitHub CEO says EU AI Act Shouldn't Apply"; Business Software Association, "BSA Leads Joint Industry Statement on the EU Artificial Intelligence Act and High-Risk Obligations for General Purpose AI."; Google Submission on EU AI Act.

Our interest here lies not in examining the pros and cons of 'open' AI itself, but in better understanding the many things 'open' AI is being used to describe, and examining the way that current arguments for and against openness in AI are being deployed in ongoing public and policy discourses–who is making which arguments, and how these arguments map to the affordances various definitions of open AI actually enable. This led us to the following conclusions:

1. At one end of a long gradient, there are a handful of maximally open AI efforts – these are non-corporate AI efforts that go to the lengths possible to offer meaningful transparency, reusability, and extensibility. But developing these models still requires access to costly computational infrastructure, which is usually leased from large tech companies.[18] In addition they require significant funding to support the development and maintenance of many complex components, including painstakingly crafted open datasets, and extensive documentation.

2. There is a long history and clear playbook for industry capture and instrumentation of open source projects, and major AI companies recognize the value of open source AI in leveraging the benefits of owning the ecosystem, enjoying the fruits of community labor, and defining the terms of engagement.

3. The ideology of the open source software movement is frequently mapped onto the concept of 'open' AI in ways that fail to account for the significant differences between large AI systems and traditional software. This shapes a narrative that assumes 'open' AI can on its own level the playing field, promote innovation, and democratize development and use. While some of this projection likely has its origin in the muddy popular definition of AI, and many contributors to open AI efforts do not espouse these views, major AI players strategically deploy such rhetorics in ways calibrated to entrench their power, often under the banner of democratized access.

3 What is (and is not) open about Open Source AI

What openness looks like in practice varies considerably, and can include anything from making the training and evaluation datasets used to shape an AI model publicly accessible, to releasing the code defining a model's architecture or its hyperparameter settings under one or another open source license, to making a model's learned parameters (model weights) publicly available, to providing documentation like model cards[19] or data sheets.[20] But some systems described as 'open' provide little more than an API and a license authorizing reuse, including commercialization of the technology.

For example, both GPTNeo and LLaMA-2 have been described as 'open' AI systems, but differ significantly in practice. While LLaMA-2 is available for free download and makes its model weights available, it fails to meet key criteria that would enable it to be conventionally considered open source.[21] Its license was drafted from scratch by Meta and isn't recognized by the Open Source Initiative, and arguably fails to meet key criteria for the Open Source Definition, which is the established definitional standard in the context of open source software.[22] It also fails to provide meaningful transparency, particularly regarding the data used to train the system. This significantly limited scope has led some to argue that LLaMA-2 does not deserve to be considered open source.[23] By contrast, GPTNeo, developed and maintained by nonprofit EleutherAI, is offered under the longstanding OSI-recognized MIT open source license, and its model weights, parameters, and detailed information about the training and configuration of the model are all provided. In addition, it was trained on an openly available dataset, the Pile[24], which was also created by EleutherAI.

Irene Solaiman's paper, which introduces the concept of 'gradients of openness', offers a useful map for thinking through the many levels of transparency, reusability, and extensibility that are possible in relation to AI's many necessary components.

[18]This remains, as a rule, true even of publicly funded infrastructures which are generally contractual arrangements with hyperscalers: AI Now Institute, "Democratize AI?: How the National AI Research Resource Falls Short." The scant exception that proves the rule is the Jean Zay supercomputer, which was built for the French government by Hewlett-Packard Enterprise and was used by BigScience to train the BLOOM model: Hewlett Packard Enterprise, "Powerful Supercomputing to Propel France's High-Performance Computing (HPC) & Artificial Intelligence (AI) Research and Development."

[19]Mitchell et al., "Model Cards for Model Reporting."

[20]Gebru et al., "Datasheets for Datasets."

[21]Vaughan-Nichols, "Meta's Llama 2 Is Not Open Source."

[22]Vaughan-Nichols.; note the OSI is undergoing a process to revise this definition to fit the particulars of open source AI development.

[23]Nolan, "Llama and ChatGPT Are Not Open-Source."

[24]Eleuther AI, "GPT Neo."

Figure 1: Source: The Gradient of Generative AI Release, Solaiman, 2023

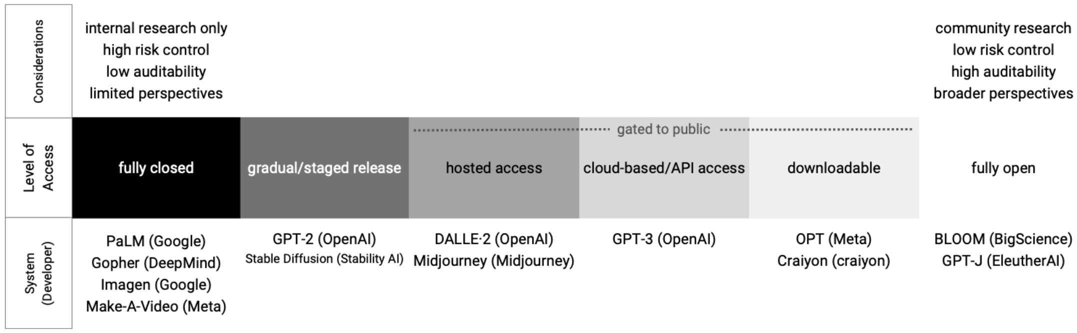

Here we review the resources–frameworks, compute, data, labor, and models–required to create and use large AI systems. This helps map a shared understanding of what comprises an AI system that lets us evaluate which parts of these systems are or can be made transparent, reusable, and extensible, which aren't or can't, and in what ways.[25]

3.1 Development frameworks

Software development frameworks make it easier for those developing software to build and deploy it in regimented, predictable, and expedient ways. They are currently part of standard development practices, creating the environment that defines how development happens, and are not unique to AI. They work by providing pre-written pieces of code, templatized workflows, evaluation tools, and other standardized methods and building blocks for common development tasks. This helps create more fungible, interoperable, and testable computational systems, while minimizing the time spent "reinventing the wheel" and avoiding bugs easily introduced when implementing systems from scratch. These components are almost always open source – this is what allows them to be used as common building blocks. They are generally reused under one or another open source license which defines the conditions under which reuse is permissible.

As with software development in general, AI development relies on a handful of popular open source development frameworks. They include increasingly vast repositories of datasets, data validation tools, evaluation tools, tools for model construction, tools for model training and export, pre-training libraries, and more, which together shape the way AI is made and deployed.

The two dominant AI development frameworks are PyTorch and TensorFlow. Both were created within large commercial technology companies Meta and Google, respectively, who continue to resource and maintain them. As of 2022, there are many more pretrained AI models that exclusively work within the PyTorch framework (including OpenAI's GPT models) than there are those that work with TensorFlow. PyTorch is also the most popular framework in academic AI research, used in the majority of academic papers.[26]

PyTorch was initially developed for internal use by Meta, but was released publicly in 2017. While PyTorch operates as a research foundation under the umbrella of the Linux Foundation, it continues to be financially supported by Meta,[27] and three of four of its lead contributors are Meta employees, while the remaining is a former Meta employee (a fairly standard practice within the tech industry when it comes to strategically important open source efforts).[28] TensorFlow was originally developed and released by Google Brain in 2015,[29] and continues to be directed and financially supported by Google, which also employs many of its core contributors.[30]

Open source development frameworks offer tools that make the AI development and deployment process quicker, more predictable, and more robust. They also have important benefits for the firms developing them. Most significantly, they

[25] Luitse and Denkena, "The Great Transformer: Examining the Role of Large Language Models in the Political Economy of AI."

[26] Foster, "PyTorch vs TensorFlow: Who Has More Pre-Trained Deep Learning Models?."

[27] Tarantola, "Meta Is Spinning off the Pytorch Framework into Its Own AI Research Foundation."

[28] Pytorch, "Contributors to Pytorch."

[29] Metz, "Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine."

[30] See Google, "Contributors to Tensorflow". Google's decision to open-source TensorFlow led to a significant change in the landscape around open source AI. While there were precursor open source frameworks including Chainer, Theano and Caffe, the quality of Google's TensorFlow package was significantly greater, and prompted Amazon, Microsoft, and Meta to release their own frameworks. See: Langenkamp and Yue, "How Open Source Machine Learning Software Shapes AI."

allow Meta, Google, and those steering framework development to standardize AI construction so it's compatible with their own company platforms – ensuring that their framework leads developers to create AI systems that, Lego-like, snap into place with their own company systems.[31] In Meta's case, this allows them to more easily integrate and commercialize academic AI models, and others developed, tuned, or deployed using PyTorch. This is true for Google and TensorFlow as well. And, in Google's case, TensorFlow has been created to easily and intuitively operate with Google's TPU hardware, the powerful proprietary computing infrastructure at the heart of Google's cloud compute business. This enables Google to optimize their commercial cloud offerings for AI development, positioning these for-profit offerings as the engine of AI development.

Open source AI development frameworks allow those bankrolling and directing them to create onramps to profitable compute offerings. They also allow Google and Meta, in this case, to shape the work practices of researchers and developers such that new AI models can be easily integrated and commercialized. This gives the company offering the framework significant indirect power within the ecosystem: training developers, researchers, and students interacting with these tools in the norms of the company's preferred framework, and thus helping define - and in some ways capture - the AI field.[32] We provide more concrete examples of this phenomenon in subsequent sections.[33]

3.2 Compute

Developing powerful AI models requires massive datasets, which require massive computational power to process.[34] Contemporary AI development is characterized by a race to scale, with datasets increasing by an order of magnitude each year and computational requirements keeping pace.[35] Access to compute presents a significant barrier to reusability for even the most maximally 'open' AI systems, because of the high cost involved in both training and running inferences on large-scale AI models at scale (i.e. instrumenting them in a product or API for widespread public use).

These significant computational requirements do not necessarily wane after the preliminary development stage, during which an AI model is initially trained and calibrated. Indeed, these upfront compute requirements can be dwarfed by the compute needed to use large AI models in the real world to provide answers or generate images.[36] Exact figures are hard to come by as this remains a closely guarded topic (hidden within contracts between infrastructure providers and businesses, as a rule), but one armchair estimate suggests that the compute required to run the ChatGPT interface exceeds the training cost of the underlying GPT-4 model on a weekly basis.[37] Another estimate from a financial analysis firm estimated Microsoft would need $4 billion in infrastructure spending to deploy ChatGPT inside its Bing search product.[38] The high cost of inference led OpenAI CEO Sam Altman to testify before congress that "we try to design systems that do not maximize for engagement. In fact, we're so short on GPUs, the less people use our products, the better".[39]

This requirement for more, and more, and more compute does not appear to be subsiding. A recently published (and then deleted) profile of OpenAI shows the company scrambling to secure more computational resources, viewing limited GPUs (specialized processors used to train AI) as the primary check on their aspiration toward bigger more powerful models.[40]

Additionally, eking out maximal computational capacity from specialized hardware requires specialized and in some cases proprietary software systems. Most AI model training is implemented in CUDA, a proprietary framework developed by Nvidia that only supports training on its proprietary GPUs.[41] OpenAI created Triton, ostensibly as an open source alternative to CUDA which may allow future portability between GPU types (platforms), but Triton currently

[31] Engler, "How Open-Source Software Shapes AI Policy."

[32] Langenkamp and Yue, "How Open Source Machine Learning Software Shapes AI.", Whittaker, "The Steep Cost of Capture."

[33] Notably, companies also often patent new AI methods at the same time that they release the attendant code openly: this protects them from intellectual property lawsuits, while facilitating mass adoption and "building on top". See for example: Shazeer et al., "Attention-based sequence transduction neural networks.", Hinton et al., "System and method for addressing overfitting in a neural network."

[34] Kaplan et al., "Scaling Laws for Neural Language Models.", Vipra and West, "Computational Power and AI."

[35] Patel, "The AI Brick Wall: A Practical Limit".

[36] While on device or 'edge' models are being used in limited applications for small-scale inference, in most instances centralized compute offers important efficiencies that matter particularly for large-scale AI.

[37] Patel, "The Inference Cost Of Search Disruption – Large Language Model Cost Analysis."

[38] Patel and Ahmad, "The Inference Cost of Search Disruption"

[39] Oremus, "ChatGPT Has Enormous Hidden Costs That Could Throttle AI Development."

[40] Habib, "OpenAI's Plans According to Sam Altman."

[41] GPU stands for Graphics Processing Units: specialized computer processors, originally developed for gaming, now primarily used for AI training because they allow many calculations to be performed quickly in parallel

only works on Nvidia's GPUs.[42] While OpenAI welcomes community contributions that could enable Triton to work on other computational platforms, it has not itself invested the resources to do so. Google's TensorFlow programs are written to enable code to be executed directly on Google's TPUs, and there are a handful of additional domain-specific compilers in TensorFlow like XLA that can accelerate the running of particular models on Google hardware. This means that while some lower level software that optimizes computational power for AI development may be open in that one can inspect their code, in practice, this software is often designed for efficiency in proprietary hardware environments, and is developed and governed by companies selling compute resources and/or licensing AI models.

In short, the computational resources needed to build new AI models and use existing ones at scale, outside of privatized enterprise contexts and individual tinkering, are scarce, extremely expensive, and concentrated in the hands of a handful of corporations, who themselves benefit from economies of scale, the capacity to control the software that optimizes compute, and the ability sell costly access to computational resources. This significant resource asymmetry undermines any claims to democratization that the availability of 'open' AI models might be used to support.

3.3 Data

Data shaped to exacting (and labor-intensive) specifications is necessary to construct large-scale AI systems. Some researchers have even claimed that access to data may be more important than access to compute when building large-scale AI.[43] Both are essential, and in the current rush to scale paradigm, the more of each, the "better" these models perform.[44]

We know very little about the data that was used to train the current generation of commercially available generative AI systems (GPT-4, Bard, PaLM-2). In particular, Google and OpenAI/Microsoft have refused to disclose important information about the training data used to shape GPT-4[45] and PaLM-2[46], not even revealing the number of tokens (discrete pieces of data) these models were trained on, let alone details on data provenance or other information necessary to understand (somewhat) the tendencies of these models.

We know a bit more about the data used to train previous generations of commercially available AI systems and most current 'open' AI models. Current 'open' AI models frequently depend on two repositories of scraped web data made publicly available for wide use: the Common Crawl dataset and the Pile. The Common Crawl dataset is a repository of scraped web page data updated monthly by a US-based non-profit of the same name, aiming to "democratize [...] data so everyone, not just big companies, can do high quality research and analysis"[47] Common Crawl is available via Amazon, whose AWS hosts the data under the company's Open Data Sponsorships program.[48] The Pile is an openly available data set designed for creating AI language systems. It was built by compiling 22 smaller datasets including Wikipedia, YouTube subtitles, and the website HackerNews. The EU-based community The Eye hosts the dataset, "in support of a persistent push toward democratizing access & research"[49] which was created by EleutherAI,[50] a non-profit focused on developing and maintaining 'open' AI. EleutherAI is supported by for-profit companies Stability AI and Hugging Face, among others, who provide funding and access to Amazon's AWS compute (through agreements between StabilityAI and Hugging Face with AWS).[51]

Scraping data to create datasets for AI development also raises issues of extraction and intellectual property. Such datasets, whether open or closed, are often assembled by taking copyrighted images, text, and code from the web, or by copying and reusing datasets compiled by language groups from the majority world like NLP Ghana and Lesan AI.[52] This means that, in practice, those using these datasets to train and evaluate AI models are often using others' work and intellectual property to do so, claiming fair use even as such claims rest on legally shaky ground.[53] Legal or not, the practice of indiscriminately trawling web data to create systems that are currently being poised to undercut the

[42] OpenAI, "Introducing Triton."
[43] Musser et al., "The Main Resource Is the Human."
[44] Kaplan et al., "Scaling Laws for Neural Language Models."
[45] OpenAI, "GPT-4 Technical Report."
[46] Elias, "Google's Newest A.I. Model Uses Nearly Five Times More Text Data for Training than Its Predecessor."
[47] "Common Crawl."
[48] "Common Crawl."
[49] Gao et al., "The Pile."
[50] Gao et al.
[51] Wiggers, "Stability AI, Hugging Face and Canva Back New AI Research Nonprofit."
[52] "Ghana Natural Language Processing (NLP).", Deck, "The AI Startup Outperforming Google Translate in Ethiopian Languages."
[53] Khan and Hanna, "The Subjects and Stages of AI Dataset Development."

livelihoods of writers, artists, and programmers—whose labor created such 'web' data in the first place—has raised alarm and ire,[54] and multiple lawsuits filed on behalf of these actors are currently moving forward.[55]

These concerns are particularly pressing when considering the colonial context in which much of this exploitation takes place. AI systems are often constructed by companies from the minority world, using data and labor resources from the majority world, and then imposed upon them without consultation.[56] As Paul Azunre, founder of the open source project Ghana NLP, recently put it, "If African AI/ML researchers are not careful, this new "Open Source" movement championed by the richest global tech companies will become a mechanism for continued exploitation of our human capital and continent".[57] The effects of these colonial dynamics compound historical patterns of inequality, reinforcing problematic stereotypes and biases.[58]

Such exploitation also runs directly counter to the notions of data sovereignty coming from the majority world, illustrated in projects such as Te Hiku Media: "While we recognize the importance of open source technology, we're mindful that the majority of *tangata whenua* and other indigenous peoples may not have access to the resources that enable them to benefit from open source technologies. . .By simply open sourcing our data and knowledge, we further allow ourselves to be colonised digitally in the modern world."[59] In keeping with Māori principles Te Hiku developed and maintains the Kaitiakitanga License, which exemplifies a practice of guardianship over data by the people from whom it originates.[60]

This is not an argument for closed datasets–closed datasets largely compound this issue. But we need to be clear about what they do, and do not, accomplish. When datasets are not made available for scrutiny, or when they are inscrutably large, it becomes very difficult to check whether these datasets launder others' intellectual property, or commercially use data that was specifically licensed for non-commercial use, or was licensed under particular sovereignty mandates. For example, Microsoft's GitHub CoPilot programming assistant—a generative AI system that produces code—has been shown to have been trained on and subsequently regurgitate code licenced under GPL,[61] an open source license that requires derivative code to be released under the same terms. However, even using permissively licensed code (eg, the MIT license) to train generative AI may similarly violate provisions requiring attribution, which current generative AI systems could, but do not provide.

Datasets like the Pile and CommonCrawl are available and licensed for wide use. But they're not ready to use off the shelf. More labor is required to make such datasets useful in the context of building large AI models. Careful construction and remixing of datasets is necessary to create performant AI: Meta's LLaMA model samples from seven different datasets in varying proportions in the process of calibrating performance.[62] BigScience's BLOOM model was trained on a composite of 498 datasets, which involved a complex data governance process as well as a manual quality filtering process to remove code, spam, and other noise.[63] The Falcon 40-B model trained by the UAE-based Technology Innovation Institute (which currently ranks best in class among open source AI models on a number of narrow evaluation tasks) foregrounds data quality in its documentation: starting from CommonCrawl, it was filtered and deduplicated, before being augmented using sources "such as research papers and conversations from social media," noting that "a particular focus was put on data quality at scale."[64] While a paper detailing the Falcon model is not yet available, a paper about its dataset reviews its data pipeline and its processes for carefully curating web data.[65]

Data labeling practices, which affix 'meaning' to the AI training and evaluation data and shape the ontology that AI systems reflect back during inference, are rarely transparent and often rely on proprietary platforms. As we review below, they are also extremely labor intensive. The most well-known companies operating in the data labeling platform space are Amazon Mechanical Turk, Sama, and ScaleAI, all of which are for-profit (as opposed to community run or academic efforts), pairing paying customers with workers paid (little) to label data on demand.

[54] FAccT, ""AI Art" and Its Impact on Artists".

[55] Deck, "The AI Startup Outperforming Google Translate in Ethiopian Languages"; Xiang, "A Photographer Tried to Get His Photos Removed from an AI Dataset. He Got an Invoice Instead," Setty, "Sarah Silverman, Authors Hit OpenAI, Meta With Copyright Suits."

[56] Browne, "AI Is Steeped in Big Tech's 'Digital Colonialism'"

[57] Paul Azunre, "If African AI/ML Researchers Are Not Careful, This New 'Open Source' Movement Championed by the Richest Global Tech Companies Will Become a Mechanism for Continued Exploitation of Our Human Capital and Continent. . . [1/3] ."

[58] Birhane, Prabhu, and Kahembwe, "Multimodal Datasets."

[59] Te Hiku Media, "He Reo Tuku Iho, He Reo Ora."

[60] "Kaitiakitanga License."

[61] Karabus, "GitHub CEO."

[62] Touvron et al., "LLaMA."

[63] Le Scao et al., "BLOOM."

[64] Technology Innovation Institute, "Falcon LLM."

[65] Penedo et al., "The RefinedWeb Dataset for Falcon LLM."

The preparation and curation of data used to train and calibrate leading large-scale AI models involves resource-intensive processes much more complicated than downloading an openly available dataset. The transparency and reusability of datasets like the Pile and CommonCrawl allow for better evaluation of model training and limitations. But beyond the cost and time required to create them in the first place, significant labor is involved in curating these before they're used in training in order to enable better model performance. In addition, as we reviewed, the cost and scarcity of compute required to process such data puts practical limits on who is able to use it in the task of creating large scale AI systems.

3.4 Labor

Large-scale AI systems' insatiable need for curated, labeled, carefully organized data means that building AI at scale requires significant human labor. This labor creates the 'intelligence' that artificial intelligence systems are marketed as automating and making computational.[66] This labor can be roughly categorized as applying to:

- Data labeling and classification
- Model calibration (reinforcement learning with human feedback, and similar processes)
- Content moderation, trust and safety and other forms of post-deployment support
- Engineering and product development and maintenance

Generative AI systems, the large-scale AI systems currently receiving the most attention, are trained and evaluated on a broad range of human-generated text, speech and/or imagery. The process of shaping a model such that it can mimic human-like output without replicating offensive or dangerous material requires intensive human involvement in order to ensure the model's outputs stay within the bounds of 'acceptable'[67]—and thus enable it to be marketed, sold, and applied in the real world by corporations and other institutions intent on maintaining customers and their reputations. This process is often called reinforcement learning with human feedback, or RLHF, which is a technical-sounding term that in practice refers to thousands of hours of often traumatic human labor. In a typical RLHF process, workers might be instructed to select which of a few snippets of text produced by a generative AI system most closely resembles human-generated text. The workers' choices would be fed back into the system, 'teaching' it what is and is not 'human-like'.[68] These processes are also applied to 'teach' a model to recognize and filter toxic, offensive, or dangerous output, work that has been repeatedly demonstrated to harm the mental health of the workers performing it.[69] This work is often outsourced, providing distance between the company developing and marketing the model and the detrimental working conditions involved in the training process.[70] Google's Bard model relies on workers recruited by outsourcing firms Accenture Plc and Appen Ltd, who task these workers with making consequential decisions about the boundaries of 'acceptable' expression even though the workers reportedly receive minimal training and work under frenzied deadlines.[71] OpenAI accomplished this for their GPT models by hiring workers in Kenya through the outsourcing firm Sama. This work resulted in harmful consequences, as workers were forced to view and read horrific ideas and images repeatedly for low wages with no meaningful support.[72] These workers have since unionized,[73] and filed a petition with the Kenyan National Assembly to investigate the welfare and working conditions of Kenyans performing such services and whether they are compliant with protections from exploitation and the right to fair remuneration and reasonable working conditions.[74]

The precarity, harm, and colonial dynamics of these labor practices raise a host of serious questions about the costs and consequences of large-scale AI development overall. And while data preparation and model calibration require this extensive, rarely heralded labor which is fundamental in creating the meaning of the data that shapes AI systems,

[66] Williams, "The Exploited Labor Behind Artificial Intelligence.", Taylor, "The Automation Charade."

[67] Thylstrup and Talat, "Detecting 'Dirt' and 'Toxicity.'"

[68] OpenAI, "Learning from Human Preferences."

[69] Workers at Meta and Microsoft have filed lawsuits for the detrimental consequences of similar labor practices in content moderation, see: Newton, "Facebook Will Pay $52 Million in Settlement with Moderators Who Developed PTSD on the Job." and Levin, "Moderators Who Had to View Child Abuse Content Sue Microsoft, Claiming PTSD."

[70] Wong, "America Already Has an AI Underclass."

[71] Alba, "Google's AI Chatbot Is Trained by Humans Who Say They're Overworked, Underpaid and Frustrated."

[72] Perrigo, "OpenAI Used Kenyan Workers on Less Than $2 Per Hour.", Hao, "The Hidden Workforce That Helped Filter Violence and Abuse Out of ChatGPT."

[73] Perrigo, "150 African Workers for AI Companies Vote to Unionize."

[74] Mercy Sumbi, "On Behalf of the Young Kenyans Whose Lives Have Been Ruined Because They Did the Dirty Work Training the #ChatGPT Algorithm, We Have Filed a Petition to @NAssemblyKE to Investigate How @OpenAI and @samasource Got Away with Such Exploitation and to Urgently Regulate Tech Work. Https://T.Co/9seeyGKqFM."

companies generally release little if any information about the labor practices underpinning this data work. Nor is failing to release such information generally criticized as a form of 'closedness'. What we do know is largely the product of either investigative journalism,[75] or organizing by researchers and workers.[76]

The labor required to curate, prepare data and calibrate systems is poorly paid, but it still costs a significant amount given the number of workers and time required to shape the data to build contemporary AI systems. This presents another barrier to democratic and 'open' access to the resources required to create and deploy large AI models (even as we cannot accept the term democratic for a structure that relies on low-paid, precarious workers who receive little benefit while enduring harm, and are themselves excluded from such imagined democracy).

3.5 AI models

A model refers to an algorithmic system that has been 'trained' and evaluated using large amounts of structured data in order to produce statistically likely outputs in response to a given input. For example, ChatGPT works by applying the GPT models, which were trained on extensive amounts of text data, much of it scraped from the web, to the task of predicting a statistically-plausible text based answer in response to a question or prompt. Once trained, an AI model can be released in the same way other software code would be released–under an open source license for reuse, leaked, or otherwise made available online. Reusing an already-trained AI model does not require having access to the underlying training or evaluation data. In this sense, many AI systems that are labeled 'open' are playing very loose with the term. Instead of providing meaningful documentation and access, they're effectively wrappers of closed models, inheriting undocumented data, failing to provide annotated RLHF training data and rarely publishing their findings - and documenting their findings in independently reviewed publications even more sparingly.[77]

There are a number of examples of large-scale 'open' AI models, available for public reuse: these include Meta's LLaMA-2; Falcon 40-B, developed by the UAE's Technology Innovation Institute, trained on AWS; MosaicML's MPT models, now tied to Microsoft's Azure, BigScience's BLOOM model; trained on the Jean Zay French supercomputer, and run on AWS; StabilityAI's Stable Diffusion; trained on AWS; and several models provided by the nonprofit LAION; which trains on Stable Diffusion's hardware, which runs on AWS. But clustering these all under the singular label of 'open' does a disservice to the serious distinctions between them.

Companies like Hugging Face and StabilityAI offer open source AI models to their customers and the public. Their business models rely not on licensing proprietary models themselves, but on charging for extra features and services on top of open models, features such as API access, model training on custom data, and security and technical support as a paid service to clients.[78] They also offer to fine tune private models for their clients, honing and calibrating the performance of already-trained models for a given task or domain.

The non-profit EleutherAI also offers large-scale open source AI models along with documentation and the codebases used to train them. EleutherAI is focused solely on fostering research on large-scale AI, licensing its models under the (very permissive) Apache 2.0 open source license for use by AI researchers. Among those engaging in 'open' AI, EleutherAI offers arguably the most maximally 'open' AI system. It operates through donations and sponsorship by CoreWeave, Hugging Face, StabilityAI, Google TRC, Nat Friedman (the former CEO of GitHub) and LambdaLabs.

A handful of academic projects have also produced large 'open' AI models at smaller scales. These include the Vicuna 'open' AI model, which was developed by a team of researchers at UCSD, CMU, and Berkeley; the Koala open AI model, which was developed at Berkeley; and Stanford's Alpaca model. The latter is well known for having been developed to run on a single laptop—a notable feat given the computationally intensive nature of deploying such models.[79] However, even a chatbot based on this extremely computationally efficient model became too costly– and risky, due to the model's hallucinations - to continue running, and the team has since taken it down.[80]

Ever-increasing model (and data, and compute) scale is the current paradigm in AI development. Bigger models are more powerful, and more resource intensive, and thus more difficult to produce outside of large technology companies. Scale can be measured in different ways, but is frequently assessed based on the number of discrete pieces of data, or tokens, in a training dataset. The largest-scale openly available AI models (LLaMA-2, MosaicML, OpenLLaMA) are trained on datasets of around 1 trillion tokens. Other well-known 'open' AI models (BLOOM, the Pythia suite of

[75] Hao and Hernandez, "How the AI Industry Profits from Catastrophe.", Perrigo, "OpenAI Used Kenyan Workers on Less Than $2 Per Hour."

[76] Perrigo, "150 African Workers for AI Companies Vote to Unionize.", Irani and Silberman, "Turkopticon."

[77] Liesenfeld, Lopez, and Dingemanse, "Opening up ChatGPT."

[78] Contrary Research, "Hugging Face.", Smith, "Leaked Deck Raises Questions over Stability AI's Series A Pitch to Investors."

[79] Edwards, "You Can Now Run a GPT-3-Level AI Model on Your Laptop, Phone, and Raspberry Pi | Ars Technica."

[80] Germain, "Stanford Researchers Take Down Alpaca AI Over Cost and Hallucinations."

models) were trained on around ∼2-300 billion tokens, comparable to OpenAI's GPT-3. None of these compare with the largest commercially available closed models: OpenAI has not released the size of its GPT-4, but Google's PaLM-2 is reportedly 3.6 trillion tokens.[81]

Much of the current action in 'open' AI is focused on the less computationally expensive task of fine-tuning models openly released by the few powerful actors who can afford to train them from scratch. Fine-tuning allows third parties to take a previously trained model and tweak it for use in a particular context or domain.[82] This is not the same as creating the capacity to build an AI model from scratch: a critical distinction, given that the process of creating from scratch involves many highly important editorial decisions that fundamentally shape the model.

The LLaMA Leak and Fine-tuning Open Source AI Models
Meta's LLaMA model was leaked to the general public after Meta shared it with a limited number of researchers, and it now forms the basis of many fine-tuning projects which are built on top of Meta's architecture, outlined below.
Many 'open' AI projects build on top of Meta's LLaMA language model,[83] which the company released as an open-source package that was intended to be given only to researchers who wrote asking for access. Unsurprisingly, one week after launch, a downloadable torrent of LLaMA was posted on 4chan and rapidly made available for wide use.[84]
LLaMA is not packaged for easy use by the casual third party. But as a pre-trained model—in which the considerable labor and expense of shaping data and conducting training has already been expended—it does reduce the compute barriers for AI developers who can use it as a starting place to build on top of.
LLaMA's leak also benefits Meta, the company that created it. A leaked memo by a Google engineer explains why: "Paradoxically, the one clear winner in all of this is Meta. Because the leaked model was theirs, they have effectively garnered an entire planet's worth of free labor. **Since most open source innovation is happening on top of their architecture, there is nothing stopping them from directly incorporating it into their products."** (emphasis added[85]
Mark Zuckerberg has said as much himself on recent Meta earnings calls, as well, which we discuss below. This is why Meta's subsequently releasing LLaMA-2 with a license enabling the model to be tweaked for commercial use makes sense from a business perspective.
Deploying LLaMA-2 (or any other large AI model) for public use — whether fine tuned or not — is also resource intensive, requiring significant computational resources that almost always require leasing cloud infrastructure from one or a handful of large tech companies.

4 A brief history of open source

The close relationship between openness and corporate power did not start with AI. Briefly accounting for the history of open source and its relationship to the tech industry provides a rough tableau against which today's complex incentives and often contradictory rhetoric of openness can be more easily understood.

The Free Software Foundation was incorporated in 1985 as part of a movement to allow people to use the new crop of personal computers without relying on proprietary, corporate-controlled software.[86] This tendency was, of course, threatening to dominant tech firms like Microsoft whose profits and growth relied on selling proprietary software. In leaked internal memos from 1998, Microsoft executives were blunt, describing open source software as a "short-term revenue and platform threat", and expressing concerns that the movement's "free idea exchange" would "present a long term developer mindshare threat".[87] To combat this perceived threat to profits and dominance, Microsoft subsequently used "everything in [their] marketing arsenal to discredit [open source software's] reliability", adopting familiar propagandistic rhetoric that described open source licenses using "three of the most feared words in the United States: cancer, communism, and un-American".[88]

The animosity of industry toward open source was not absolute, however. Some companies figured out how to benefit from openly developed and freely available software, and with that efforts to align business incentives to open source practices gained traction. The Open Source Initiative was founded in 1998, adopting the "open source" label as a way to distance themselves from what they saw as the "moralizing and confrontational attitude that had been associated with free software, and to instead make the case for open software principles on "pragmatic, business-case grounds".[89]

[81]Elias, "Google's Newest A.I. Model Uses Nearly Five Times More Text Data for Training than Its Predecessor."
[82]Hu et al., "LoRA: Low-Rank Adaptation of Large Language Models."
[86]Williams, "Free as in Freedom."
[87]Microsoft, "Open Source Software: A (New?) Development Methodology"
[88]Coleman, "Coding Freedom."
[89]Open Source Initiative, "History of the OSI.", see also: Open Source Initiative, "Open Source Case for Business."

These business-case grounds become clearer in a contemporaneous example. In 1999, IBM invested $1B in the open source operating system Linux. They did this not so much out of enlightened good will, but in large part as a play to interrupt Microsoft's market dominance. MarketWatch described IBM's investment as an attempt to "steal market share from proprietary Unix flavors, as well as Microsoft's Windows".[90] Here, among other things, corporate investment in open source acted as a way to weaken competitors.

We see a similar dynamic when we look at Google and the Android operating system. In 2005, Google acquired Android, releasing Android OS as open source software in 2007. This move gave Google a number of advantages in the mobile market, and positioned them as an alternative to Apple. By releasing Android OS openly, and providing ample development frameworks, incentives, and support, Google was able to attract interest from developers who devoted their time to creating and maintaining applications for Android. These, in turn, made the operating system more attractive to consumers who wanted to use these apps, driving adoption. Notably, some enforcement agencies have made clear that being open source doesn't disqualify an entity from engaging in anti-competitive behavior: the UK Competition and Markets Authority's (CMA) 2022 mobile ecosystem market study concluded that there is effectively a duopoly in the mobile market. While Google's decision to offer the Android operating system under an open source license encouraged adoption and competition at the application layer, it did not reduce barriers to entry in operating system development, and switching between the Android and iOS operating systems remains challenging.[91] The CMA plans to take further investigative action,[92] and the European Commission has also recently concluded a market investigation into Google Android for anticompetitive practices[93] by issuing the company a $4.3 billion fine for "restricting mobile competition and consumer choice".[94]

The close relationship between corporate power and open source software is also evidenced in the case of Amazon and MongoDB. In 2019, Amazon implemented its own version of the popular open source database MongoDB, and released it to consumers as DocumentDB, a proprietary service sold as part of Amazon's AWS cloud offering.[95] This was after the leader the MongoDB company criticized Amazon and other cloud companies for "capturing all the value [of open source] and giving nothing back to the community"[96] — value created by volunteer contributors and reviewers who helped build and test MongoDB as a critical resource they and others relied on (even as MongoDB was operated as a for-profit company). By forking this open source project, hosting it on their own proprietary infrastructure, and whitelabeling it as its own, Amazon was able to exploit and integrate the work of open source developers into its own service offerings.

Similar logic undergirds corporate strategies around 'open' AI, as well. During a 2023 earnings call, Meta CEO Mark Zuckerberg explicitly laid out the company's rationale for releasing its dominant PyTorch AI framework openly:[97]

> "[PyTorch] has generally become the standard in the industry as a tool that a lot of folks who are – who are building AI models and different things in that space use, it's generally been very valuable for us [. . .] because it's integrated with our technology stack, when there are opportunities to make integrations with products, it's much easier to make sure that developers and other folks are compatible with the things that we need in the way that our systems work."[98]

PyTorch's library of easily accessible development tools functions as a de-facto standard, and in turn allows Meta to easily integrate new models and other innovations into their products for profit, in many cases without paying for or directly managing their development.

Taken together, these examples show that where some tech companies initially fought open source, seeing it as a threat to their own proprietary offerings, more recently these companies have tended to embrace it as a mechanism that can allow them to entrench dominance by setting standards of development while benefiting from the free labor of open source contributors. Throughout open source history, we see that while an "open regime can lead to increased innovation [and] economic activity . . . in the end the major benefits are being harvested by a small group of companies".[99]

[90] Marketwatch, "Market Dynamics: The Battle for Enterprise Linux."

[91] Competition and Markets Authority, "Mobile Ecosystems Market Study Final Report."

[92] Competition and Markets Authority, "Mobile Ecosystems Market Study.", Competition and Markets Authority, "Mobile Browsers and Cloud Gaming"

[93] European Commission, "Antitrust: Commission opens formal investigation against Google in relation to Android mobile operating system."

[94] Chan, "Google Appeals Huge Android Antitrust Fine to EU's Top Court."

[95] Kanaracus, "AWS, MongoDB Database Collision Stirs Open Source Tensions."

[96] Claburn, "Fed up with Cloud Giants Ripping off Its Database, MongoDB Forks New 'Open-Source License.'"

[97] Foster, "PyTorch vs TensorFlow: Who Has More Pre-Trained Deep Learning Models?."

[98] Meta, "Q1 Earnings Call Transcript".

[99] McCann, "Commoning Intellectual Property: Public Funding and the Creation of a Knowledge Commons."

This is not to argue against open source, or to deny its many benefits, which are often mixed in complex ways with dynamics of corporate capture and influence. But in order to view the relationship between 'open' AI and corporate concentration through a clear lens, we need to recognize that 'open' is not inherently democratizing, especially in an ecosystem increasingly defined by concentrated capital and infrastructural power.

The many roads to to capture and corporate benefit: a few brief case studies
Companies have captured or used open source projects to their benefit in a variety of ways:

1. **Invest in open source to attack your proprietary competitors: IBM and Linux.**
   In 1999, IBM invested $1B in the open source operating system Linux to curb Microsoft's market dominance.
2. **Release open source to control a platform: Google and Android.**
   In 2007, Google open sourced and heavily invested in Android OS, allowing them to achieve mobile operating system prominence versus competitor Apple, and attracting scrutiny from regulators for anticompetitive practices.
3. **Re-implement and sell as SAAS: Amazon and MongoDB.**
   In 2019, Amazon implemented its own version of the popular open source database MongoDB, and then sold it as a service on its AWS platform.
4. **Develop an open source framework that enables the firm to integrate open source products into its proprietary systems: Meta and PyTorch.**
   Meta CEO Mark Zuckerberg has described how open sourcing the PyTorch framework has made it easier to capitalize on new ideas developed externally and for free.

5 The case of OpenAI, LP: openness as marketing

In 2015, OpenAI launched as a nonprofit research lab with a mission to "build value for everyone rather than for shareholders"[100] and a stated policy that encouraged all employees to "publish their work, whether as papers, blog posts, or code."[101] The company also announced that it would share any patents it filed for with the world at large. In other words, at the outset, OpenAI promoted a model of "openness" that aspired to share intellectual property, research findings, and some financial information via mandatory nonprofit tax filings.

Even at this early stage, this is a limited version of 'openness'. It does not include open governance, democratic or collaborative contributions, open data, or openly available access to the costly resources required to create large scale AI models. OpenAI's founders–with the exception of Elon Musk, an early funder[102]–have also consistently voiced ambivalence about the risks associated with releasing its tech widely: "Doing all your research in the open is not necessarily the best way to go . . . We will produce a lot of open source code. But we will also have a lot of stuff that we are not quite ready to release", Greg Brockman, OpenAI's Chair, told WIRED in 2016.[103]

Their ambivalence was exemplified by the release of GPT-2 in 2019. In February 2019, the company announced it had created GPT-2, large language model, but described it as "too dangerous"[104], and chose not to release it out of concern that the system could be misused "to generate deceptive, biased, or abusive language at scale".[105] Such rhetoric also

[100] OpenAI, "Introducing OpenAI."

[101] Initial projects included OpenAI Gym, a toolkit for building and comparing reinforcement learning algorithms, and Universe, which enables any program to be turned into a Gym environment.

[102] Elon Musk, "@GRDecter OpenAI Was Created as an Open Source (Which Is Why I Named It 'Open' AI), Non-Profit Company to Serve as a Counterweight to Google, but Now It Has Become a Closed Source, Maximum-Profit Company Effectively Controlled by Microsoft. Not What I Intended at All."

[103] Metz, "Inside OpenAI, Elon Musk's Wild Plan to Set Artificial Intelligence Free."

[104] The company did release a smaller version of the model at the time, indicating it associated scale with safety. OpenAI, "We've Trained an Unsupervised Language Model That Can Generate Coherent Paragraphs and Perform Rudimentary Reading Comprehension, Machine Translation, Question Answering, and Summarization — All without Task-Specific Training: Https://Blog.Openai.Com/Better-Language-Models/ Https://T.Co/360bGgoea3."

[105] OpenAI, "Better Language Models and Their Implications." Note that OpenAI was not the only entity at work building such systems: as then-lead of Google's Ethical Artificial Intelligence Team Margaret Mitchell said in response to the announcement, "Let's please not pretend there are not awesome people who have been dealing with the potential issues of inevitable progress, ahead of time, for years. The relevant experts are not limited to the people who did the thing. The relevant experts include people who foresaw this obvious sequence and chose not to do the thing". Google had several similar systems in development, as Mitchell went on to discuss with her coauthors in the Stochastic Parrots paper.

acts as marketing, painting these models as capable and powerful (even if the danger of such power is highlighted). And they present contradictions when we look at the reality of the AI business model that Microsoft and Open AI are pursuing: here Brockmon paints these models as too capable to be openly licensed. But, as evidenced by their business offerings, clearly Microsoft and Open AI don't view them as too capable to be sold via Microsoft Azure API contracts, from which they are currently available for third party use.

OpenAI's plan for a controlled release quickly failed: by June 2019, graduate student Connor Leahy had managed to replicate GPT-2, though after a series of discussions with OpenAI and researchers at the Machine Intelligence Research Institute (which studies "existential risk" in AI), Leahy opted not to release it, saying "this isn't just about GPT2. What matters is that at some point in the future, someone *will* create something truly dangerous and there need to be commonly accepted safety norms *before* that happens."[106] In August, another group of researchers at Brown University replicated and opted to publish the model, emphasizing that "because of the relative ease of replicating this model, an overwhelming number of interested parties could replicate GPT-2 [. . .] Because our replication efforts are not unique [. . .] we believe releasing our model is a reasonable first step towards countering the potential future abuse of these kinds of models".[107] OpenAI then published its own version of the model in November 2019, claiming it saw "no strong evidence" of misuse.[108]

OpenAI co-founder Ilya Sutskever has since voiced regret about this choice: "We were wrong [...] I fully expect that in a few years it's going to be completely obvious to everyone that open-sourcing AI is just not wise".[109] Events in the intervening years may have had much to do with the deepening of their reticence to openness: in March 2019 OpenAI announced that it was abandoning its non-profit status, transforming into OpenAI LP, a "capped-profit" company, meaning that first round investors cannot receive returns in excess of 100 times their initial investment, which in the context of the billions of dollars of investment in OpenAI, LP, is not a meaningful limitation, which is why we reference OpenAI as for profit throughout this paper. OpenAI justified this change by saying it would allow the company to "rapidly increase our investment in compute and talent while including checks and balances to actualize our mission."[110] In July 2019, OpenAI announced it would receive a $1 billion investment from Microsoft,[111] exclusively licensing its GPT-3 model to the company, and effectively folding Open AI into Microsoft as a division of the corporation.[112]

Along with this change in OpenAI's organizational structure came a transition from limited and aspirational openness to all but total closed-ness. When OpenAI published its GPT-4 model, following the company's integration into Microsoft, the organization kept fundamental details about the system obscured from the public view. GPT-4's technical report explicitly declined to release details about GPT-4's "architecture (including model size), hardware, training compute, dataset construction, [and] training method".[113] Some of these details were kept secret not just out of concern for their safety implications, but due to OpenAI's desire to remain ahead of competitors — something that clearly shows the transition from their initial imperative of building "value for everyone rather than for shareholders."

6 The arguments For and Against 'Open' AI, and Who's Making Them

After exploring what is (and is not) "open" about current 'open' AI systems, alongside a brief history of open source and free software, we now broaden our view to look at how the rhetoric of openness is being deployed in current AI policy debates.

Openness is currently being referenced by powerful actors in policy discussions in an effort to shape the trajectory of regulatory policy around AI. Conversations and understandings about 'open' AI have significant and high-stakes implications beyond the narrow terms of open source licenses, transparency mechanisms, or extensibility, and are currently being deployed to shape the AI policy landscape overall. As with any examination of lobbying and influence, the rhetoric around 'open' AI policy advocacy must be read in light of the particular interests of the entity making them. And as in the past, openness is today being used by companies as a rhetorical wand to lobby to entrench and bolster their positions.

[106]Notably, the cloud computing cost to train a model of the size of GPT 1.5 was $40k, see [R] OpenAI: Better Language Models and Their Implications. Leahy said in one of his posts that "I feel like I owe Google my firstborn child or something for the amount of free support they"ve given me".
[107]Cohen, "OpenGPT-2: We Replicated GPT-2 Because You Can Too.."
[108]OpenAI, "GPT-2."
[109]Vincent, "OpenAI Co-Founder on Company's Past Approach to Openly Sharing Research."
[110]OpenAI, "OpenAI LP."
[111]Microsoft, "OpenAI Forms Exclusive Computing Partnership with Microsoft to Build New Azure AI Supercomputing Technologies."
[112]Hao, "OpenAI Is Giving Microsoft Exclusive Access to Its GPT-3 Language Model."
[113]OpenAI, "GPT-4 Technical Report."

The debate around whether or not to exempt 'open source AI' from regulations mandated under the EU AI Act offers an instructive example. In the lead-up to the European Parliament's vote on the draft text, a cluster of organizations expressed concern that the EU's proposed regulation would be too onerous for open source AI developers. This argument was led by the non-profit organization LAION which issued an open letter claiming regulatory intervention would stifle innovation.[114] The requirements they rebuked included mandates to conduct risk assessments for systems deployed in 'high risk' contexts, and to ensure adequate documentation and traceability of automated decision making produced by AI models (which arguably fit within the "transparency and reusability" definition of openness to begin with).[115]

While LAION played a significant role, the most active lobbying came from industry. In Google's submission to the EU regarding the AI Act, the company claimed that the Act would have "a chilling effect on open collaboration in the AI ecosystem".[116] The CEO of Microsoft's GitHub argued that open source AI was creating a "spring of innovation across the world and here in Europe",[117] while the Business Software Alliance - an industry trade group that represents Cisco, IBM, and Oracle, among others - argued that compliance with the Act "would severely impact and disincentivize the development of open source software and AI in Europe".[118]

Why were industry players so active in pushing for open source AI to be excluded from the mandates of the EU's AI Act?

As illustrated above, large, entrenched tech companies like Google, Microsoft, Meta and others have vested interests in 'open' AI development. The author of the leaked 'Google Moat' memo said the quiet part out loud: "the value of owning the ecosystem" to the company "cannot be overstated": such ecosystem capture directly contributed to Google's dominance across the domains of search, mobile, and advertising.[119] Similarly, as the platform where much open source code is hosted (code that has been scraped and used to train models like Microsoft's CoPilot), Microsoft's Github has its own vested interests in unimpeded open source development, even as Microsoft-supported company OpenAI remains much more circumspect as it pushes for licensing (enabling Microsoft to tacitly support contradictory standpoints that both, in the end, benefit the company). To protect their interests, many companies are lobbying for exemptions to baseline documentation and accountability mechanisms.

With this big picture in mind, we move now to examine the claims underlying lobbying arguments being made about open source AI:

6.1 "Open AI creates safety through transparency"

LAION's open letter argues that open-source AI promotes safety by enabling researchers and authorities to audit model performance, identify risks, and establish mitigations or countermeasures.[120] We agree with this claim in general. However, the resources and access needed to conduct such audits and tests are almost exactly those that 'open' AI projects would have been required to produce and conduct by the requirements in the AI Act, making the opposition to including such models under the AI Act confusing, assuming safety was the cardinal goal.

The effectiveness of auditing as a safety measure is heavily predicated on ensuring that significant resources are available and incentives aligned such that meaningful audits actually take place, and are robust enough to account for the real risks posed by the deployment of AI models. Simply making the code open does not guarantee that expert volunteers will spend their time and resources thoroughly reviewing it. One need only look to the Heartbleed vulnerability as an example showing that openly scrutable code does not necessarily mean that such code is audited and free from errors — even when such code is as critically important as OpenSSL, the encryption protocol that effectively enabled commerce on the internet.[121] The Heartbleed vulnerability, if exploited, could enable information transmitted over the web - like peoples' passwords and other private information - to be exfiltrated. The vulnerability was not detected earlier due in part to insufficient resourcing of audits of the open source code: OpenSSL was operating on a budget of $2,000 a year

[114] Wiggers, "A Startup Wants to Democratize the Tech behind DALL-E 2, Consequences Be Damned."; LAION itself is supported in part by Hugging Face and Stability AI, two for-profit companies whose business model revolves around 'open' AI.

[115] European Commission, "Regulatory Framework Proposal on Artificial Intelligence | Shaping Europe's Digital Future."

[116] Google, "Submission to the European Commission on the EU AI Act."

[117] Karabus, "GitHub CEO Says EU AI Act Shouldn't Apply to Open Source AI."

[118] https://www.bsa.org/news-events/news/bsa-leads-joint-industry-statement-on-the-eu-artificial-intelligence-act-and-high-risk-obligations-for-general-purpose-ai "BSA Leads Joint Industry Statement on the EU Artificial Intelligence Act and High-Risk Obligations for General Purpose AI | BSA | The Software Alliance."

[119] Patel, "Google 'We Have No Moat, And Neither Does OpenAI.'"

[120] Laion, "A Call to Protect Open-Source AI in Europe."

[121] Synopsys, "Heartbleed Bug."

in donations despite being used to secure two-thirds of the world's websites.[122] Following the revelation of the bug Amazon, Google, and Microsoft, among others, formed a "Core Infrastructure Initiative" to which they each donated $300,000 to maintain the OpenSSL code,[123] still a scant pittance compared to the resources dedicated to corporate development efforts within these companies.

What Heartbleed demonstrates is that just because code *can* be audited does not mean that it *will* be. Technically 'open' code and documentation is not itself sufficient to ensure careful review and remediation: this dynamic is exacerbated in the case of AI, where large probabilistic models function in ways that produce outputs and decisions that cannot be predicted from the code, documentation, and data alone. Thus, audits that account for the impact and implications of AI systems deployed in complex domains require significantly more resources and access than just opening up code and READMEs. This makes it less likely that such audits will happen without clear and established incentives and processes. Similarly, arguments that open source AI will lead to safer AI must also speak to how auditing, maintenance, and inspection will occur, not simply to the fact that the various transparency attributes of open source AI *could* enable such inspection, as well as where the responsibilities lies when vulnerabilities and flaws are left unaddressed.

6.2 "Open AI increases insecurity"

Arguments in the opposing direction position open source AI as a source of deep *insecurity*, by making powerful technology widely available for reuse, potentially placing it in the hands of bad actors. As the security researchers Bruce Schneier and Jim Waldo put it (while still arguing in favor of open source AI), "having the technology open-sourced means that those who wish to use it for unintended, illegal, or nefarious purposes have the same access to the technology as anyone else."[124]

This fear purportedly motivated OpenAI in its initial decision not to release GPT-2 as open source, and continues to be reflected in statements from OpenAI leaders like Ilya Sutskever, who said "These models are very potent and they're becoming more and more potent. At some point it will be quite easy, if one wanted, to cause a great deal of harm with those models. And as the capabilities get higher it makes sense that you don't want to disclose them."[125]

Despite this, in May 2023, we saw OpenAI advocate against "burdensome mechanisms like licenses or audits" for today's corporate and open source AI projects which they say "create tremendous value in the world" (such as their GPT-4), even as the company continued to advocate for such regulatory action for AI with capability significantly beyond today's levels. This wields regulation as a moat, advocating against regulation for their dominant model, but for regulation for anything significantly more capable which may challenge it.

It's worth noting here that we've described two instances of lobbying through two entities tightly tied to Microsoft, seemingly in different directions. GitHub's argument, outlined in section one, is self interested because they (and Microsoft) rely on open source development, both as a business model for the GitHub platform and as a source of training data for profitable systems like CoPilot. This makes sense for them, while OpenAI is arguing primarily that models "above a certain threshold" should not be open — a threshold that they effectively set due to resource monopolies Microsoft benefits from. So, open source exceptions are good for them. Arguing that open sourcing their powerful models is dangerous also benefits OpenAI — this claim both reasserts the power of their models and allows them to conflate resource concentration with cutting edge scientific development.

Concerns about safety were prominent in Senators' Blumenthal and Hawley's letter to Meta questioning the company about its decision making prior to releasing the first LLaMA model, which evidenced significant skepticism about Meta's decision to release the model: "Meta's choice to distribute LLaMA in such an unrestrained and permissive manner raises important and complicated questions about when and how it is appropriate to openly release sophisticated AI models. Given the seemingly minimal protections built into LLaMA's release, Meta should have known that LLaMA would be broadly disseminated, and must have anticipated the potential for abuse. While Meta has described the release as a leak, its chief AI scientist has stated that open models are key to its commercial success".[126]

Concerns about insecurity emanating from the proliferation of open source AI models are justified, among other reasons because open source models enable AI to be fine-tuned at *small* scale without a steep learning curve. What remains unremarked and unclear about this argument, especially given those making it, is why access to the same or similarly powerful models to those obtained through a cloud contract from Microsoft or Google—which is the current standard—poses less danger than reusing an openly released AI model. Further, assuming that openly released AI

[122] Perlroth, "Heartbleed Highlights a Contradiction in the Web."
[123] Finkle, "Big Tech Companies Offer Millions after Heartbleed Crisis | Reuters."
[124] Schneier and Waldo, "Big Tech Isn't Prepared for A.I.'s Next Chapter."
[125] Vincent, "OpenAI Co-Founder on Company's Past Approach to Openly Sharing Research."
[126] Blumenthal, "Letter to Meta."

models are in fact more dangerous, this does not mean that the countervailing benefits of a more concentrated market offer the best path forward: industry concentration is creating toxic competition among AI firms, leading them to release models commercially before they are ready and before they have undergone necessary scrutiny or risk mitigation.[127] A small number of large players does not, of its own accord, ensure safer AI.

6.3 "Open AI will reduce industry concentration"

Another set of claims posits that regulation will introduce requirements that are far too onerous for open source developers to comply with, thus contributing to concentration in AI. For example, one proposal that Sam Altman put forth in his testimony before Congress would be to create some kind of licensing regime for generative AI systems. Different to the kinds of permissible use licensing that typically accompany open source software development, such an approach would likely instantiate a gatekeeping entity that would ensure some level of scrutiny of systems or companies before AI models are made available for commercial use.

In a recent article Sayash Kapoor and Arvind Narayanan argued that requiring AI providers to obtain licenses of this kind would worsen AI risks by fostering a monoculture and outcome homogenization - essentially that having such a licensing regime would mean that only certain types of models that fit the regime are developed. In contrast, they position open source AI as a viable alternative to ensure greater diversity in AI development.[128] While reliance on a few corporate-produced models is indeed a pressing concern, homogenization characterizes the already existing AI environment: homogenization is a product of the existing high levels of concentration and capture of the academic field of AI (and the infrastructure needed to build it),[129] rather than a future outcome of proposed regulatory interventions. And as outlined above, open AI exists within that environment, not outside of it: as recent research into open source AI vulnerabilities illustrates, the proliferation of a small handful of open source AI models would introduce the same problems of a single point of failure as heavy dependency on a small handful or closed models produced by large tech firms.[130]

The LAION letter makes a modified version of Kapoor and Narayan's argument, claiming that open source AI will enable entry by small businesses which can build on existing models through fine-tuning rather than simply purchasing access to models created by large firms.[131] What this claim ignores is the many critical elements baked into the originating large scale models and open source libraries: the fine-tuned end products largely function as barnacles on the hull of Big Tech, rather than a meaningful alternative to it. They still need to be run on Big Tech infrastructures (as a rule), and cede power to define and create the core model logics to the large companies who have the resources to create them from scratch.

The computer scientist Zeerak Talat's refutation of the LAION letter further points out that open source exemptions could in fact *increase*, rather than reduce, the burden on small businesses. Exempting open source models from regulatory scrutiny could create an undue competitive burden for smaller firms, by ensuring that compliance requirements fall on the entity that fine-tunes an end product rather than the larger actor that developed the original large-scale model or library.[132]

6.4 "Open AI is key to innovation"

A final set of arguments advocates for exemptions to regulation on the basis that open source AI is critical to innovation. For example, Betsy Masiello and Derek Slater argue that "open source AI is driving innovation in powerful AI models", while caveating this claim by noting that innovation will come in the form of commoditizing products built off large language models (most of which are created by large tech firms themselves, or by efforts using large tech firms" cloud infrastructures).[133] In a more circumspect statement Alex Engler has outlined that open source AI can speed AI adoption by reducing the level of technical knowledge necessary to implement AI.[134] Schneier and Waldo similarly claim that "Rather than needing tens of thousands of machines and millions of dollars to train a new model, an existing model can now be *customized* on a mid-priced laptop in a few hours".[135]

[127] "Computational Power and AI - AI Now Institute.", Blumenthal, "Letter to Meta."
[128] Kapoor, "Licensing Is Neither Feasible nor Effective for Addressing AI Risks."
[129] Whittaker, "The Steep Cost of Capture."
[130] Gu, Dolan-Gavitt, and Garg, "BadNets."
[131] "A Call to Protect Open-Source AI in Europe."
[132] Talat, "Response to the LAION Letter."
[133] Slater, "Will Open Source AI Shift Power from 'Big Tech'?"
[134] Langenkamp and Yue, "How Open Source Machine Learning Software Shapes AI."
[135] Schneier and Waldo, "Big Tech Isn't Prepared for A.I.'s Next Chapter."

These statements are true. But they also offer a narrowly defined interpretation of innovation - one that does not meaningfully trouble the incentive structures, business models, or the dynamics that dictate who is a ‘user’ of AI systems vs whom such systems are ‘used on’. Nor does it meaningfully shift who benefits from the development of AI systems, and who risks the harms.

Rather than offer a true alternative to AI controlled by a handful of large corporations, the primary locus of open source AI development and ‘innovation’ is in the fine tuning of models developed by large technology firms, which in turn require ongoing compute contracts to run inference (i.e. to use in real life) at any significant scale. Meanwhile, the focus on fine tuning corporate-trained models, and extending these to various domains, also functions as free product development for these same companies, who as we have seen have a long history of exploiting and capturing the most successful open source projects.[136] This dynamic is compounded by the current corporate control and direction setting of key AI development frameworks, which as we’ve reviewed work to propagate critical elements in AI construction in ways that ensure the resulting models and implementations will be compatible with large companies’ proprietary systems.

Other claims around open source-driven innovation are even less nuanced. The claims made by Thibault Schrepel and Alex Pentland in a recent article, notably that ‘open source and open access’ AI be exempted from antitrust scrutiny, provide a clear example: “If open source and open access players were able to share up-front costs, marketing networks, and technical knowledge," they argue, “They would be in a stronger position to compete with proprietary systems that capture more of the value they create."[137] The conflation of ‘open source and open access’ introduced here would include any AI model being offered through an API - a framing encompassing almost the entire market of large language models.

Finally, as the analysis above elucidates, ‘openness’ often enables systemic exploitation of developers’ and creators’ labor while maintaining the infrastructural and ecosystem dominance of the largest firms.[138] In the context of high levels of corporate concentration and gatekeeping over the ingredients necessary to build AI systems, ‘open’ AI as currently operationalized is primed for corporate capture.

7 Conclusion

Even in its more maximal instantiations, in which ‘open’ AI systems provide robust transparency, reusability, and extensibility, such affordances do not, on their own, ensure democratic access to or meaningful competition in AI. Nor does openness alone solve the problem of AI oversight and scrutiny. Even so, the rhetoric and promise of openness in AI systems is being leveraged by powerful companies to bolster their positions in the face of growing interest in AI regulation. With some companies embracing ‘open’ AI as a mechanism to entrench dominance, investing in ‘open’ AI efforts in ways that allow them to set standards of development while benefiting from the free labor of open source contributors.

Policymakers need to approach the task of regulating AI with a clear understanding of the many things AI is, and is not, and with a materially grounded recognition of what ‘open’ AI can, and cannot, deliver. This will produce a vastly different picture of the affordances of ‘open’ AI than that being painted in much of the current rhetoric. It will also require focusing on the significant differences between open source software and ‘open’ AI, and recognizing that the development processes, resource requirements, and inherent centralization of AI mean that it cannot be easily described or defined in terms forged originally to promote and define open source software. And just as AI is distinct from traditional software, so too are the early 2020s distinct from the halcyon days of the late 1990s, when open source software cemented its ideological moorings amid a rush to commercialized networked computation.

While concerns about safety and concentration should be addressed, whether through or in spite of open source AI, they should be considered as part of a much larger picture. If anything, ramping up already heavy levels of market concentration itself creates safety harms, such as the introduction of single points of failure through which systemic risk may be disseminated.[139] And there are many further harms that go far beyond safety considerations to take into account regarding the deployment of large predictive systems into sensitive social and political domains.

Similarly ‘openness’ in AI does not *directly* lead to innovation, despite claims to the contrary; certainly not when significant industry presence, corporate ownership of the costly infrastructure needed to make new AI systems, and

[136] See, for example, Amazon’s forking MongoDB and whitelabeling it as DocumentDB.

[137] Schrepel and Pentland, “Competition between AI Foundation Models."

[138] See, for example: https://twitter.com/pazunre/status/1569743778524680192; https://www.vice.com/en/article/pkapb7/a-photographer-tried-to-get-his-photos-removed-from-an-ai-dataset-he-got-an-invoice-instead

[139] https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/03/inquiry-cloud-computing-business-practices-federal-trade-commission-seeking-public-comments

corporate capture of the field of academic AI research[140] forms the backdrop for 'open' AI research. Evidence indicates these factors are already contributing to the development of an AI monoculture,[141] incentivizing the creation of only the kinds of AI systems that industry is most invested in and able to profit from.[142]

Creating meaningful alternatives to technology defined and dominated by large, monopolistic corporations is and remains a laudable goal - one that requires creating alternatives to corporate-dominated infrastructure. This is a particularly pressing need as AI systems are integrated into many highly sensitive domains with particular public impact: in healthcare, finance, education, and the workplace, AI has diffuse and profound effects that should not be determined by a small handful of profit-motivated companies, and that cannot be understood simply from examining system code and documentation in a vacuum. The creation of meaningful alternatives won't be accomplished through pursuit of even maximally 'open' AI development alone. We need a wider scope for AI development and greater diversity of methods, as well as the construction of AI and other systems that more meaningfully attend to the needs of the public, not of commercial interests. Creating the conditions to make such alternatives possible is a project that can coexist with, and even be supported, by, regulation.[143] But pinning our hopes on open source AI in isolation won't lead us to that world, and in many respects could make things far worse.

8 Acknowledgements

Our analysis builds on the work of past scholars who have written on connections between open source software, AI, and corporate power: Nick Dyer-Witherford, Atle Mikkola Kjosen and James Steinhoff's exploration of how open source software interacts with commercial logics, Seda Gürses, Jennifer Cobbe, Jatinder Singh and Michael Veale's work foregrounding AI and computational infrastructures, Gabriella Coleman, Chris Kelty, and Dawn Nafus's scholarship on open source developer communities, and the work of Irene Solaiman and Pablo Chavez challenging binary views of software systems as uniformly 'open' or 'closed'.

We hope that others build on their work and this to further illuminate relationships between corporate power, open source and AI. We are grateful to Paul Azunre, Amanda Bertsch, Stella Biderman, Francois Chollet, Sireesh Gururaja, Alex Hanna, Amba Kak, Chris Kelty, Josh Lund, Keoni Mahelona, Varoon Mathur, Margaret Mitchell, Aviya Skowron, Jed Sundwall, Zeerak Talat, and Luis Villa for feedback and conversations that deeply enriched this paper.

9 References

"Paul Azunre on Twitter: 'If African AI/ML Researchers Are Not Careful, This New "Open Source" Movement Championed by the Richest Global Tech Companies Will Become a Mechanism for Continued Exploitation of Our Human Capital and Continent. . . [1/3]' / X." Accessed August 10, 2023. https://twitter.com/pazunre/status/1569743778524680192.

Ahmed, Nur, Muntasir Wahed and Neil C. Thompson. "The Growing Influence of Industry in AI Research." Accessed August 10, 2023. https://www.science.org/doi/10.1126/science.ade2420.

Amadeo, Ron. "Chrome's New Ad-Blocker-Limiting Extension Platform Will Launch in 2023." Accessed August 10, 2023. https://arstechnica.com/gadgets/2022/09/chromes-new-ad-blocker-limiting-extension-platform-will-launch-in-2023/.

Birhane, Abeba, Vinay Uday Prabhu, and Emmanuel Kahembwe. "Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes." arXiv, October 5, 2021. https://doi.org/10.48550/arXiv.2110.01963.

Bloomberg Professional Services. "AI Presents Nearly Level Playing Field for Startups and Big Tech." March 21, 2023, Bloomberg. https://www.bloomberg.com/professional/blog/ai-presents-nearly-level-playing-field-for-startups-and-big-tech/.

Bloomberg.com. "Google's AI Chatbot Is Trained by Humans Who Say They're Overworked, Underpaid and Frustrated." July 12, 2023. https://www.bloomberg.com/news/articles/2023-07-12/google-s-ai-chatbot-is-trained-by-humans-who-say-they-re-overworked-underpaid-and-frustrated.

[140] Whittaker, "The Steep Cost of Capture."
[141] Ahmed et al, "The Growing Influence of Industry in AI Research."
[142] Hooker, "The Hardware Lottery."
[143] See, for example, the EPA's investments in funding participatory science projects which include grant-making to community and citizen science organizations: US EPA, "Funding Opportunities for Participatory Science Projects."

Browne, Grace. “AI Is Steeped in Big Tech’s ‘Digital Colonialism’." Accessed August 10, 2023. https://www.wired.com/story/abeba-birhane-ai-datasets/.

Business Software Alliance. “BSA Leads Joint Industry Statement on the EU Artificial Intelligence Act and High-Risk Obligations for General Purpose AI." Accessed August 10, 2023. https://www.bsa.org/news-events/news/bsa-leads-joint-industry-statement-on-the-eu-artificial-intelligence-act-and-high-risk-obligations-for-general-purpose-ai.

Chan, Kelvin. “Google Appeals Huge Android Antitrust Fine to EU’s Top Court," December 1, 2022. https://apnews.com/article/technology-business-european-union-commission-operating-systems-f5633c74eedea662582025d909fe5fb6.

Claburn, Thomas. “Fed up with Cloud Giants Ripping off Its Database, MongoDB Forks New ‘Open-Source License.’" Accessed August 10, 2023. https://www.theregister.com/2018/10/16/mongodb_licensning_change/.

Clegg, Nick. “Nick Clegg: Openness on AI Is the Way Forward for Tech." Financial Times, July 11, 2023.

Cohen, Vanya. “OpenGPT-2: We Replicated GPT-2 Because You Can Too." Medium (blog), December 11, 2019. https://medium.com/@vanya_cohen/opengpt-2-we-replicated-gpt-2-because-you-can-too-45e34e6d36dc.

Coldewey, Devin. “OpenAI Shifts from Nonprofit to ‘capped-Profit’ to Attract Capital." TechCrunch (blog), March 11, 2019. https://techcrunch.com/2019/03/11/openai-shifts-from-nonprofit-to-capped-profit-to-attract-capital/.

Coleman, E. Gabriella. “Coding Freedom: The Ethics and Aesthetics of Hacking." In Coding Freedom. Princeton University Press, 2012. https://doi.org/10.1515/9781400845293.

Common Crawl. “Frequently Asked Questions." Accessed August 10, 2023. https://commoncrawl.org/big-picture/frequently-asked-questions/.

Commonwealth. “Commoning Intellectual Property: Public Funding and the Creation of a Knowledge Commons." Accessed August 10, 2023. https://www.common-wealth.org/publications/commoning-intellectual-property.

Competition and Markets Authority. “Mobile Ecosystems Market Study Final Report," August 4, 2022. https://www.gov.uk/government/publications/mobile-ecosystems-market-study-final-report.

Competition and Markets Authority. “Mobile Ecosystems Market Study," June 10, 2022. https://www.gov.uk/cma-cases/mobile-ecosystems-market-study.

Contrary. “Contrary Research." Accessed August 10, 2023. https://contraryresearch.substack.com/embed.

Crunchbase. “OpenAI - Funding, Financials, Valuation & Investors." Accessed August 10, 2023. https://www.crunchbase.com/organization/openai/company_financials.

Deck, Andrew.. “The AI Startup Outperforming Google Translate in Ethiopian Languages," July 11, 2023. https://restofworld.org/2023/3-minutes-with-asmelash-teka-hadgu/.

Doffman, Zak. “Why You Should Delete Google Chrome After New Tracking Admission." Forbes. Accessed August 10, 2023. https://www.forbes.com/sites/zakdoffman/2021/08/28/stop-using-google-chrome-on-windows-10-android-and-apple-iphones-ipads-and-macs/.

Edwards, Benj. “You Can Now Run a GPT-3-Level AI Model on Your Laptop, Phone, and Raspberry Pi." Accessed August 10, 2023. https://arstechnica.com/information-technology/2023/03/you-can-now-run-a-gpt-3-level-ai-model-on-your-laptop-phone-and-raspberry-pi/.

Eleuther AI. “GPT Neo." Accessed August 10, 2023. https://huggingface.co/docs/transformers/model_doc/gpt_neo.

Elias, Jennifer. “Google’s Newest A.I. Model Uses Nearly Five Times More Text Data for Training than Its Predecessor." CNBC, May 16, 2023. https://www.cnbc.com/2023/05/16/googles-palm-2-uses-nearly-five-times-more-text-data-than-predecessor.html.

Elon Musk [@elonmusk]. “@GRDecter OpenAI Was Created as an Open Source (Which Is Why I Named It ‘Open’ AI), Non-Profit Company to Serve as a Counterweight to Google, but Now It Has Become a Closed Source, Maximum-Profit Company Effectively Controlled by Microsoft. Not What I Intended at All." Tweet. Twitter, February 17, 2023. https://twitter.com/elonmusk/status/1626516035863212034.

Engler, Alex. “How Open-Source Software Shapes AI Policy." Accessed August 10, 2023. https://www.brookings.edu/articles/how-open-source-software-shapes-ai-policy/.

Escalante-DeMattei, Shanti. “Artists Sue Midjourney, Stability AI — The Case Could Change Art." Accessed August 10, 2023. https://www.artnews.com/art-in-america/features/midjourney-ai-art-image-generators-lawsuit-1234665579/.

European Commission. “Antitrust: Commission investigates Google’s Android practices." Text. Accessed August 10, 2023. https://ec.europa.eu/commission/presscorner/detail/it/MEMO_15_4782.

European Commission. “European Commission - Have Your Say." Text. Accessed August 10, 2023. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/F2662492_en.

European Commission. “Regulatory Framework Proposal on Artificial Intelligence," June 30, 2023. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai.

FAccT. “ACM FAccT - 2023 Accepted CRAFT Sessions." Accessed August 10, 2023. https://facctconference.org/2023/acceptedcraft.html.

Federal Trade Commission. “An Inquiry into Cloud Computing Business Practices: The Federal Trade Commission Is Seeking Public Comments," March 21, 2023. https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/03/inquiry-cloud-computing-business-practices-federal-trade-commission-seeking-public-comments.

Finkle, Jim. “Big Tech Companies Offer Millions after Heartbleed Crisis." Accessed August 10, 2023. https://www.reuters.com/article/uk-cybercrime-heartbleed-idINKBN0DA0SE20140424.

Foster, Kelsey. “PyTorch vs TensorFlow: Who Has More Pre-Trained Deep Learning Models?." Accessed August 10, 2023. https://hackernoon.com/pytorch-vs-tensorflow-who-has-more-pre-trained-deep-learning-models.

Gao, Leo, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, et al. “The Pile: An 800GB Dataset of Diverse Text for Language Modeling." arXiv, December 31, 2020. http://arxiv.org/abs/2101.00027.

Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. “Datasheets for Datasets." arXiv, December 1, 2021. https://doi.org/10.48550/arXiv.1803.09010.

Germain, Thomas. “Stanford Researchers Take Down Alpaca AI Over Cost and Hallucinations," March 21, 2023. https://gizmodo.com/stanford-ai-alpaca-llama-facebook-taken-down-chatgpt-1850247570.

Ghana Natural Language Processing (NLP). “Ghana Natural Language Processing (NLP) - Ghana NLP." Accessed August 10, 2023. https://ghananlp.org/.

GitHub. “Contributors to Pytorch/Pytorch." Accessed August 10, 2023. https://github.com/pytorch/pytorch.

GitHub. “Contributors to Tensorflow/Tensorflow." Accessed August 10, 2023. https://github.com/tensorflow/tensorflow.

GitHub. “Kaitiakitanga License." Accessed August 10, 2023. https://github.com/TeHikuMedia/Kaitiakitanga-License/blob/tumu/LICENSE.md.

Goeva, Aleksandrina, Sara Stoudt, and Ana Trisovic. “Toward Reproducible and Extensible Research: From Values to Action." Harvard Data Science Review 2, no. 4 (July 15, 2020). https://doi.org/10.1162/99608f92.1cc3d72a.

Google. “Submission to European Commission on the EU AI Act." Accessed August 10, 2023. https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12527-Artificial-intelligence-ethical-and-legal-requirements/F2662492_en.

Gu, Tianyu, Brendan Dolan-Gavitt, and Siddharth Garg. “BadNets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain." arXiv, March 11, 2019. https://doi.org/10.48550/arXiv.1708.06733.

gwern. “Thanks. So Then It W. . .." Reddit Comment. R/MachineLearning, February 15, 2019. www.reddit.com/r/MachineLearning/comments/aqlzde/r_openai_better_language_models_and_their/eghxt8n/.

Habib, Raza. “OpenAI’s Plans According to Sam Altman." Accessed August 10, 2023. https://web.archive.org/web/20230601000258/https://website-nm4keew22-humanloopml.vercel.app/blog/openai-plans.

Hao, Karen and Andrea Paola Hernandez. “How the AI Industry Profits from Catastrophe | MIT Technology Review." Accessed August 10, 2023. https://www.technologyreview.com/2022/04/20/1050392/ai-industry-appen-scale-data-labels/.

Hao, Karen. “OpenAI Is Giving Microsoft Exclusive Access to Its GPT-3 Language Model." Accessed August 10, 2023. https://www.technologyreview.com/2020/09/23/1008729/openai-is-giving-microsoft-exclusive-access-to-its-gpt-3-language-model/.

Hao, Karen. “The Hidden Workforce That Helped Filter Violence and Abuse Out of ChatGPT - The Journal." Accessed August 10, 2023. https://www.wsj.com/podcasts/the-journal/the-hidden-workforce-that-helped-filter-violence-and-abuse-out-of-chatgpt/ffc2427f-bdd8-47b7-9a4b-27e7267cf413.

Hewlett-Packard Enterprise. “Powerful Supercomputing to Propel France’s High-Performance Computing (HPC) & Artificial Intelligence (AI) Research and Development." Accessed August 10, 2023. https://www.hpe.com/us/en/newsroom/blog-post/2019/01/powerful-supercomputing-to-propel-frances-high-performance-computing-hpc-artificial-intelligence-ai-research-and-development.html.

Hinton, Geoffrey E., Alexander Krizhevsky, Ilya Sutskever, and Nitish Srivastva. System and method for addressing overfitting in a neural network. United States US9406017B2, filed August 30, 2013, and issued August 2, 2016. https://patents.google.com/patent/US9406017B2/en.

Hooker, Sara. “The Hardware Lottery." arXiv, September 21, 2020. https://doi.org/10.48550/arXiv.2009.06489.

Hu, Edward J., Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. “LoRA: Low-Rank Adaptation of Large Language Models." arXiv, October 16, 2021. https://doi.org/10.48550/arXiv.2106.09685.

Kak, Amba, Brittany Smith, Meredith Whittaker and Sarah Myers West. “Democratize AI? How the Proposed National AI Research Resource Falls Short." AI Now Institute, October 5, 2021. https://ainowinstitute.org/publication/democratize-ai-how-the-proposed-national-ai-research-resource-falls-short.

Kanaracus, Chris. “AWS, MongoDB Database Collision Stirs Open Source Tensions | TechTarget." Accessed August 10, 2023. https://www.techtarget.com/searchaws/news/252455751/AWS-MongoDB-database-collision-stirs-open-source-tensions.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. “Scaling Laws for Neural Language Models." arXiv, January 22, 2020. https://doi.org/10.48550/arXiv.2001.08361.

Kapoor, Sayash and Arvind Narayanan. “Licensing Is Neither Feasible nor Effective for Addressing AI Risks," March 20, 2023. https://www.aisnakeoil.com/p/licensing-is-neither-feasible-nor.

Karabus, Jude. “GitHub CEO: EU AI Act Shouldn’t Apply to Open Source Devs." Accessed August 10, 2023. https://www.theregister.com/2023/02/07/github_ceo_ai_act/.

Khan, Mehtab, and Alex Hanna. “The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability." SSRN Scholarly Paper. Rochester, NY, September 13, 2022. https://doi.org/10.2139/ssrn.4217148.

LAION. “A Call to Protect Open-Source AI in Europe." Accessed August 10, 2023. https://laion.ai/notes/letter-to-the-eu-parliament.

Langenkamp, Max, and Daniel N. Yue. “How Open Source Machine Learning Software Shapes AI." In Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, 385–95. Oxford United Kingdom: ACM, 2022. https://doi.org/10.1145/3514094.3534167.

Leswing, Kif. “Meet the $10,000 Nvidia Chip Powering the Race for A.I." CNBC, February 23, 2023. https://www.cnbc.com/2023/02/23/nvidias-a100-is-the-10000-chip-powering-the-race-for-ai-.html

Levin, Sam. “Moderators Who Had to View Child Abuse Content Sue Microsoft, Claiming PTSD." Accessed August 10, 2023. https://www.theguardian.com/technology/2017/jan/11/microsoft-employees-child-abuse-lawsuit-ptsd.

Liesenfeld, Andreas, Alianda Lopez, and Mark Dingemanse. “Opening up ChatGPT: Tracking Openness, Transparency, and Accountability in Instruction-Tuned Text Generators." In Proceedings of the 5th International Conference on Conversational User Interfaces, 1–6, 2023. https://doi.org/10.1145/3571884.3604316

Lipton, Zachary C. “The Mythos of Model Interpretability." arXiv, March 6, 2017. https://doi.org/10.48550/arXiv.1606.03490.

Luitse, Dieuwertje and Wiebke Denkena, “The Great Transformer: Examining the Role of Large Language Models in the Political Economy of AI." Accessed August 10, 2023. https://journals.sagepub.com/doi/full/10.1177/20539517211047734.

Maffuli, Stephano. “Towards a Definition of ‘Open Artificial Intelligence’: First Meeting Recap." Voices of Open Source (blog), July 13, 2023. https://blog.opensource.org/towards-a-definition-of-open-artificial-intelligence-first-meeting-recap/.

Maffulli, Stefano. “Meta’s LLaMa 2 License Is Not Open Source." Voices of Open Source (blog), July 20, 2023. https://blog.opensource.org/metas-llama-2-license-is-not-open-source/.

Mercy Sumbi [@MercyMutemi]. “On Behalf of the Young Kenyans Whose Lives Have Been Ruined Because They Did the Dirty Work Training the #ChatGPT Algorithm, We Have Filed a Petition to @NAssemblyKE to Investigate How @OpenAI and @samasource Got Away with Such Exploitation and to Urgently Regulate Tech Work. Https://T.Co/9seeyGKqFM." Tweet. Twitter, July 12, 2023. https://twitter.com/MercyMutemi/status/1678984336996028416.

Metz, Cade. “Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine | WIRED." Accessed August 10, 2023. https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/.

Metz, Cade. “Inside OpenAI, Elon Musk’s Wild Plan to Set Artificial Intelligence Free." Accessed August 10, 2023. https://www.wired.com/2016/04/openai-elon-musk-sam-altman-plan-to-set-artificial-intelligence-free/.

Microsoft. “Open Source Software: A (New?) Development Methodology." Accessed August 10, 2023. http://www.catb.org/~esr/halloween/halloween1.html.

Microsoft. “OpenAI Forms Exclusive Computing Partnership with Microsoft to Build New Azure AI Supercomputing Technologies." Stories, July 22, 2019. https://news.microsoft.com/2019/07/22/openai-forms-exclusive-computing-partnership-with-microsoft-to-build-new-azure-ai-supercomputing-technologies/.

Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. “Model Cards for Model Reporting." In Proceedings of the Conference on Fairness, Accountability, and Transparency, 220–29, 2019. https://doi.org/10.1145/3287560.3287596.

MMitchell [@mmitchell_ai]. “It’s Really Hard to Watch the GPT-2 Conversations Unfold like so Much Else in Tech. 1/." Tweet. Twitter, February 18, 2019. https://twitter.com/mmitchell_ai/status/1097626427048964098.

Musser, Micah, Rebecca Gelles, Ronnie Kinoshita, Catherine Aiken And Andrew Lohn. “‘The Main Resource Is the Human’." Accessed August 10, 2023. https://cset.georgetown.edu/publication/the-main-resource-is-the-human/.

Newton, Casey. “Facebook Will Pay $52 Million in Settlement with Moderators Who Developed PTSD on the Job." The Verge, May 12, 2020. https://www.theverge.com/2020/5/12/21255870/facebook-content-moderator-settlement-scola-ptsd-mental-health.

Nolan, Michael. “Llama and ChatGPT Are Not Open-Source." Accessed August 10, 2023. https://spectrum.ieee.org/open-source-llm-not-open.

Open Source Initiative. “History of the OSI." Accessed August 10, 2023. https://web.archive.org/web/20111127235213/http://www.opensource.org/history.

Open Source Initiative. “Open Source Case for Business:Advocacy," March 31, 2007. https://opensource.org/case_for_business-php/.

OpenAI [@OpenAI]. “We’ve Trained an Unsupervised Language Model That Can Generate Coherent Paragraphs and Perform Rudimentary Reading Comprehension, Machine Translation, Question Answering, and Summarization — All without Task-Specific Training: Https://Blog.Openai.Com/Better-Language-Models/ Https://T.Co/360bGgoea3." Tweet. Twitter, February 14, 2019. https://twitter.com/OpenAI/status/1096092704709070851.

OpenAI. “Better Language Models and Their Implications." Accessed August 10, 2023. https://web.archive.org/web/20190215000029/https://blog.openai.com/better-language-models/.

OpenAI. “GPT-2: 1.5B Release." Accessed August 10, 2023. https://openai.com/research/gpt-2-1-5b-release.

OpenAI. “GPT-4 Technical Report." arXiv, March 27, 2023. https://doi.org/10.48550/arXiv.2303.08774.

OpenAI. “Introducing OpenAI." Accessed August 10, 2023. https://openai.com/blog/introducing-openai.

OpenAI. “Introducing Triton: Open-Source GPU Programming for Neural Networks." Accessed August 10, 2023. https://openai.com/research/triton.

OpenAI. “Learning from Human Preferences." Accessed August 10, 2023. https://openai.com/research/learning-from-human-preferences.

OpenAI. “OpenAI Gym Beta." Accessed August 10, 2023. https://openai.com/research/openai-gym-beta?ref=hackernoon.com.

OpenAI. “OpenAI LP." Accessed August 10, 2023. https://openai.com/blog/openai-lp.

OpenAI. “Universe." Accessed August 10, 2023. https://openai.com/research/universe.

OpenLM Research. “OpenLLaMA: An Open Reproduction of LLaMA." 2023. Reprint, OpenLM Research, August 10, 2023. https://github.com/openlm-research/open_llama.

Oremus, Will. “ChatGPT Has Enormous Hidden Costs That Could Throttle AI Development." Accessed August 10, 2023. https://www.washingtonpost.com/technology/2023/06/05/chatgpt-hidden-cost-gpu-compute/.

Patel, Dylan and Shazeer Ahmed. “The Inference Cost Of Search Disruption – Large Language Model Cost Analysis." Accessed August 10, 2023. https://www.semianalysis.com/p/the-inference-cost-of-search-disruption.

Patel, Dylan. “Google ‘We Have No Moat, And Neither Does OpenAI.’" Accessed August 10, 2023. https://www.semianalysis.com/p/google-we-have-no-moat-and-neither.

Penedo, Guilherme, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. “The RefinedWeb Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only." arXiv, June 1, 2023. https://doi.org/10.48550/arXiv.2306.01116.

Perlroth, Nicole. “Heartbleed Highlights a Contradiction in the Web." Accessed August 10, 2023. https://www.nytimes.com/2014/04/19/technology/heartbleed-highlights-a-contradiction-in-the-web.html.

Perrigo, Billy. “150 African Workers for AI Companies Vote to Unionize." Accessed August 10, 2023. https://time.com/6275995/chatgpt-facebook-african-workers-union/.

Perrigo, Billy. “OpenAI Used Kenyan Workers on Less Than $2 Per Hour: Exclusive." Accessed August 10, 2023. https://time.com/6247678/openai-chatgpt-kenya-workers/.

Poursabzi-Sangdeh, Forough, Daniel G. Goldstein, Jake M. Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. “Manipulating and Measuring Model Interpretability." arXiv, August 15, 2021. https://doi.org/10.48550/arXiv.1802.07810.

RedPajama. “RedPajama-Data: An Open Source Recipe to Reproduce LLaMA Training Dataset." Python. 2023. Reprint, Together, August 10, 2023. https://github.com/togethercomputer/RedPajama-Data.

Rudin, Cynthia. “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead." Nature Machine Intelligence 1, no. 5 (May 2019): 206–15. https://doi.org/10.1038/s42256-019-0048-x.

Schneier, Bruce, and Jim Waldo. “Big Tech Isn’t Prepared for A.I.’s Next Chapter." Slate, May 30, 2023. https://slate.com/technology/2023/05/ai-regulation-open-source-meta.html.

Schrepel, Thibault, and Alex Pentland. “Competition between AI Foundation Models: Dynamics and Policy Recommendations." SSRN Scholarly Paper. Rochester, NY, June 28, 2023. https://papers.ssrn.com/abstract=4493900.

Setty, Ridhi. “Sarah Silverman, Authors Hit OpenAI, Meta With Copyright Suits." Accessed August 10, 2023. https://news.bloomberglaw.com/ip-law/sarah-silverman-authors-hit-openai-meta-with-copyright-suits.

Shazeer, Noam M., Aidan Nicholas Gomez, Lukasz Mieczyslaw Kaiser, Jakob D. Uszkoreit, Llion Owen Jones, Niki J. Parmar, Illia Polosukhin, and Ashish Teku Vaswani. Attention-based sequence transduction neural networks. United States US10452978B2, filed June 28, 2018, and issued October 22, 2019. https://patents.google.com/patent/US10452978B2/en.

Slater, Derek. "Will Open Source AI Shift Power from 'Big Tech'? It Depends." Tech Policy Press, June 16, 2023. https://techpolicy.press/will-open-source-ai-shift-power-from-big-tech-it-depends/.

Smith, Tim. "Leaked Deck Raises Questions over Stability AI's Series A Pitch to Investors." Accessed August 10, 2023. https://sifted.eu/articles/stability-ai-fundraise-leak.

Solaiman, Irene. "The Gradient of Generative AI Release: Methods and Considerations." arXiv, February 5, 2023. https://doi.org/10.48550/arXiv.2302.04844.

sonic0002. "GitHub Copilot May Generate Code Containing GPL Code." Accessed August 10, 2023. https://www.pixelstech.net/article/1682104779-GitHub-Copilot-may-generate-code-containing-GPL-code.

Substack. "Https://Substack.Com/Channel-Frame." Accessed August 10, 2023. https://substack.com/channel-frame.

Suchman, Lucy. "What Is AI? Part 2, with Lucy Suchman." Accessed August 10, 2023. https://ainowinstitute.org/general/what-is-ai-part-2-with-lucy-suchman-ai-now-salons.

Surman, Mark. "Introducing Mozilla.Ai: Investing in Trustworthy AI." Accessed August 10, 2023. https://blog.mozilla.org/en/mozilla/introducing-mozilla-ai-investing-in-trustworthy-ai/.

Synopsys. "Heartbleed Bug." Accessed August 10, 2023. https://heartbleed.com/.

Talat, Zeerak. "Response to the LAION Letter." Medium (blog), May 17, 2023. https://medium.com/@zeeraktalat/response-to-the-laion-letter-96ed68b2936.

Tarantola, Andrew. "Meta Is Spinning off the Pytorch Framework into Its Own AI Research Foundation." Accessed August 10, 2023. https://www.engadget.com/meta-is-spinning-off-its-pytorch-ai-framework-into-its-own-research-foundation-140051987.html.

Taylor, Astra. "The Automation Charade." Logic(s) Magazine. Accessed August 10, 2023. https://logicmag.io/failure/the-automation-charade/.

Te Hiku Media. "He Reo Tuku Iho, He Reo Ora: Living Language Transmitted Intergenerationally." MAI Journal: A New Zealand Journal of Indigenous Scholarship 11, no. 1 (November 1, 2022): 41–49. https://doi.org/10.20507/MAIJournal.2022.11.1.4.

Technology Innovation Institute. "Falcon LLM." Accessed August 10, 2023. https://falconllm.tii.ae/.

The Free Library. "Market Dynamics: The Battle for Enterprise Linux.." Accessed August 10, 2023. https://www.thefreelibrary.com/Market+dynamics%3A+the+battle+for+enterprise+Linux.-a0105204889.

Thomson, Martin. "A Privacy Analysis of Google's Topics Proposal," n.d.

Thylstrup, Nanna, and Zeerak Talat. "Detecting 'Dirt' and 'Toxicity': Rethinking Content Moderation as Pollution Behaviour." SSRN Scholarly Paper. Rochester, NY, October 12, 2020. https://doi.org/10.2139/ssrn.3709719.

Touvron, Hugo, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, et al. "LLaMA: Open and Efficient Foundation Language Models." arXiv, February 27, 2023. https://doi.org/10.48550/arXiv.2302.13971.

Turkopticon. "Turkopticon." Accessed August 10, 2023. https://dl.acm.org/doi/pdf/10.1145/2470654.2470742.

US EPA, REG 02. "Funding Opportunities for Participatory Science Projects." Overviews and Factsheets, August 28, 2018. https://www.epa.gov/participatory-science/funding-opportunities-participatory-science-projects.

Vaughan-Nichols, Steven J. "Meta's Llama 2 Is Not Open Source." Accessed August 10, 2023. https://www.theregister.com/2023/07/21/llama_is_not_open_source/.

Vincent, James. . "OpenAI Co-Founder on Company's Past Approach to Openly Sharing Research: 'We Were Wrong.'" The Verge, March 15, 2023. https://www.theverge.com/2023/3/15/23640180/openai-gpt-4-launch-closed-research-ilya-sutskever-interview.

Vincent, James. “Meta’s Powerful AI Language Model Has Leaked Online — What Happens Now?" The Verge, March 8, 2023. https://www.theverge.com/2023/3/8/23629362/meta-ai-language-model-llama-leak-online-misuse.

Vipra, Jai and West, Sarah. “Computational Power and AI." Accessed August 10, 2023. https://ainowinstitute.org/publication/policy/computational-power-and-ai.

Whittaker, Meredith. “The Steep Cost of Capture by Meredith Whittaker." Accessed August 10, 2023. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4135581.

Wiggers, Kyle. “A Startup Wants to Democratize the Tech behind DALL-E 2, Consequences Be Damned." TechCrunch (blog), August 12, 2022. https://techcrunch.com/2022/08/12/a-startup-wants-to-democratize-the-tech-behind-dall-e-2-consequences-be-damned/.

Wiggers, Kyle. “Stability AI, Hugging Face and Canva Back New AI Research Nonprofit." TechCrunch (blog), March 2, 2023. https://techcrunch.com/2023/03/02/stability-ai-hugging-face-and-canva-back-new-ai-research-nonprofit/.

Williams, Adrienne. “The Exploited Labor Behind Artificial Intelligence," October 13, 2022. https://www.noemamag.com/the-exploited-labor-behind-artificial-intelligence.

Williams, Sam. “Free as in Freedom." Accessed August 10, 2023. https://www.oreilly.com/openbook/freedom/ch01.html.

Wong, Matteo. “America Already Has an AI Underclass." The Atlantic, July 26, 2023. https://www.theatlantic.com/technology/archive/2023/07/ai-chatbot-human-evaluator-feedback/674805/.

Workshop, BigScience, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, et al. “BLOOM: A 176B-Parameter Open-Access Multilingual Language Model." arXiv, June 27, 2023. https://doi.org/10.48550/arXiv.2211.05100.

Xiang, Chloe. “A Photographer Tried to Get His Photos Removed from an AI Dataset. He Got an Invoice Instead." Accessed August 10, 2023. https://www.vice.com/en/article/pkapb7/a-photographer-tried-to-get-his-photos-removed-from-an-ai-dataset-he-got-an-invoice-instead.