Why 'open' AI systems are actually closed, and why this matters

https://doi.org/10.1038/s41586-024-08141-1

Received: 5 December 2023

Accepted: 1 October 2024

Published online: 27 November 2024

Check for updates

David Gray Widder^{1⊠}, Meredith Whittaker^{2,3} & Sarah Myers West⁴

This paper examines 'open' artificial intelligence (AI). Claims about 'open' Al often lack precision, frequently eliding scrutiny of substantial industry concentration in large-scale AI development and deployment, and often incorrectly applying understandings of 'open' imported from free and open-source software to AI systems. At present, powerful actors are seeking to shape policy using claims that 'open' Al is either beneficial to innovation and democracy, on the one hand, or detrimental to safety, on the other. When policy is being shaped, definitions matter. To add clarity to this debate, we examine the basis for claims of openness in AI, and offer a material analysis of what Al is and what 'openness' in Al can and cannot provide: examining models, data, labour, frameworks, and computational power. We highlight three main affordances of 'open' AI, namely transparency, reusability, and extensibility, and we observe that maximally 'open' AI allows some forms of oversight and experimentation on top of existing models. However, we find that openness alone does not perturb the concentration of power in AI. lust as many traditional open-source software projects were co-opted in various ways by large technology companies, we show how rhetoric around 'open' AI is frequently wielded in ways that exacerbate rather than reduce concentration of power in the AI sector.

This Perspective examines 'open' artificial intelligence (AI). We find that concepts from open-source software are being applied in ill-fitting ways to AI systems. At a time when industry players are seeking to influence policy using claims that open AI is beneficial to scientific innovation and democracy, on the one hand, or detrimental to safety, on the other, we look to grounding discussions about the affordances of 'openness' in AI in a material analysis of what AI is and what openness in AI can and cannot provide.

With this aim, we review the core components of AI systems, examining which of these can be made open or not, as well as reviewing the ecosystem that has formed around the concept of open AI. We find that open AI systems can offer transparency, reusability and extensibility: they can be scrutinized, reused and built 'on top of', to varying degrees. But we also find that claims posited around openness often lack precision, frequently focused on only one stage in the development-to-deployment life cycle of AI systems, often neglecting substantial industry concentration in large-scale AI development and deployment and thus warping common-sense understandings of openness imparted from free and open-source software. Discourses that index on openness in isolation from the economic incentives of AI rarely engage issues of context, power and use-how will such systems be used, by whom, on whom-even as these issues profoundly shape the policy outcomes that debates around openness and AI claim to concern themselves with.

These questions are particularly important in our present AI landscape, which is dominated by corporate actors¹⁻⁵. Creating the conditions under which independent alternatives to industry-dominated tech can thrive is a worthy cause. However, just as many traditional open-source projects were co-opted in various ways by large technology companies, our findings indicate that the rhetoric of openness is frequently wielded in ways that, far from alleviating, instead exacerbate the concentration of power in the AI sector.

The rhetoric of open AI is at present directing political and research attention and shaping policy in both the USA and the European Union, among other jurisdictions⁶⁻⁹. The 'open-source AI' debate has been substantially constructed by AI companies, who have used claims around openness to serve their particular regulatory and market aims. Depending on their business model, companies have used the rhetoric of openness to implicitly back arguments that AI should either be exempt from regulation¹⁰ or be subject to stringent licensing requirements or export controls¹¹. Meanwhile, recent work by researchers has helpfully complicated these claims, even if it hasn't reshaped the public debate, adding nuance and grounding by evaluating the risks and benefits of model openness^{12,13} and creating taxonomies of more or less open models in an attempt to provide conceptual clarity^{14,15}.

Open AI and definitional arbitrage

The definition of AI itself is contested and unclear, further muddling the question of what 'open' means in the context of AI. Over its more than 70-year history, the term AI has been applied to a wide variety of approaches, less as a technical term of art and more as marketing and aspiration^{4,16}. Some AI systems are deterministic, such as rule-based systems, which–given a set of inputs–follow a set of instructions to

¹Digital Life Initiative, Cornell University, New York City, NY, USA. ²Signal Foundation, Mountain View, CA, USA. ³University of Western Australia, Perth, Western Australia, Australia, Australia. ⁴AI Now Institute, New York City, NY, USA. ⁵²e-mail: david.g.widder@gmail.com

produce clearly defined outputs. Others are probabilistic, making comparisons to vast pools of data and drawing inferences from the connections between data points. At present, the term often describes probabilistic, large, resource-intensive machine-learning systems, with so-called 'generative' AI attracting the most attention in popular discourse. Because large and generative AI systems most clearly perturb traditional definitions of open source and because they are the focus of present policy and discourse, we focus on these systems.

The need for definitional clarity has prompted considerable debate¹⁷, and has culminated in a proposal from the Open Source Initiative¹⁸. In more popularized discussions about AI, drawing on ideologies about free software that were forged decades ago with the aim of resisting corporate control¹⁹, conventional understandings of free and open-source software are being projected onto open AI systems even when they do not fit^{19,20}. From the promise that open source democratizes software development, that many eyes on open code could ensure its integrity and security²¹, or that open source levels the playing field and allows the innovative to triumph^{22,23}, open-source software did many of these things, to varying degrees¹⁸.

Methods of asserting dominance through—not in spite of—open-source software

Over the history of free and open-source software, for-profit tech companies have used their resources to capture ecosystems, or have used open-source projects to assert dominance in a variety of ways. Here are examples used by companies in the past.

 Invest in open source to challenge your proprietary competitors. 	IBM and Linux. In 1999, IBM invested US\$1 billion in the open-source operating system Linux—operating software positioned as an open-source alternative to the then-dominant Microsoft—and established the Linux Foundation ²⁴ .
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 prominence over competitor Apple and attracting scrutiny from regulators for anticompetitive practices ²⁵ .
3. Re-implement and sell	Amazon and MongoDB. In 2019, Amazon
as Software As A Service (SAAS).	implemented its own version of the popular open-source database MongoDB, known as DocumentDB ²⁶ , and sold it as a service on its AWS platform. In 2022, it transitioned to a revenue-sharing agreement with MongoDB ²⁷⁻²⁹ .

Open AI is a different story from open-source software in key respects. Unlike open-source software, identifying harms and flaws in AI systems requires much more than open weights and an accessible application programming interface (API) or openly licensed AI model (as in Meta's LLaMA model series), and although provision of the training data and rigorous open documentation have salutary effects on the ability to audit AI systems critical for accountability, there are inherent limitations in the ability to predict the behaviour of systems that are probabilistic³².

Likewise, although openness can foster competition at the edges– enabling others to build on top of base AI models through fine-tuning with a high level of efficiency–this does not perturb the characteristics of the market at large. Nor does fine-tuning eliminate the impact of key decisions made during the development phase of the base model³³. Factors that make the playing field in AI uneven include network effects, access to datasets, access to and cost of the computing needed for inference at scale, lack of a viable business model and, at present, inflated interest rates^{34–38}. Together, these factors strongly limit the competitiveness of AI start-ups in the present business environment and contribute to a market in which the paths to profit, by and large, channel through large tech companies—whose infrastructures are imperative for AI development and whose access to markets are imperative for any return on investment³⁹. Openness may enable greater ability to modify AI models that have already been developed but these larger environmental factors influence whether the product of such experimentation has a path to market⁴⁰.

In practice, gradients of AI openness offer greatly differing affordances, even though they are all confusingly clustered under the same term, 'openness⁴¹. Some systems described as open, such as Meta's LLaMA-3 (ref. 42), offer little more than an API or the ability to download a model subject to distinctly non-open use restrictions^{42,43}. In these cases, this is 'openwashing' systems that are better understood as closed⁴⁴⁻⁴⁶. Other maximal variants of open AI, such as EleutherAI's Pythia series, go much further, offering access to the source code, underlying training data and full documentation, as well as licensing the AI model for wide reuse under terms aligned with the Open Source Initiative's long-standing definition of open source.

Given these confused definitions, unless quoting verbatim claims, we avoid the term open source in the rest of this paper and instead use the blanket term open.

What is (and is not) open about open AI?

AI systems require distinct development processes and rely on specialized and costly resources concentrated in the hands of a few large tech companies^{5,47-49}. Given the resources required to produce large-scale Al systems, commercial Al companies with computing power, datasets and research teams have increasingly dominated the field of AI research and development. As such, these companies not only shape the trajectory of what gets built but also the conditions under which AI systems can be built, including what elements of a system (weights and datasets) are made open for others to access and reuse. Although new techniques have made it easier to build leaner, more efficient use cases that are fine-tuned with larger base models⁵⁰, they have done so without changing these underlying characteristics of the market. Ultimately, the cost and resources needed for training, and the choke points that large companies hold in terms of access to market, mean that open AI thus does not straightforwardly equate to a shift in competitive conditions for the AI market, although in its more maximal instantiations, it provides three key affordances:

- 1. **Transparency.** Many AI systems labelled 'open' publish weights, documentation or data about the system. Maximal examples of open AI provide access to the underlying training data and information about the weights associated with a given model. Both of these are useful for enabling some forms of validation and auditing^{51,52} and both help augur post-hoc insights into system behaviour that are critical for accountability. Because of the probabilistic nature of present AI systems, assertions about the transparency of AI systems, however open they are, should be measured, particularly when drawing comparisons with traditional software: knowing the weights, code and documentation cannot tell us exactly how a model will perform in a given context or explain why a given outcome occurs or enable us to predict the so-called 'emergent' properties of the system³³⁻⁵⁵.
- 2. **Reusability.** Some open AI models and data are licensed and made available to third parties to reuse⁵⁶. Openly licensed data and model weights, and the frequent use of traditional open-source licences in making these available, have contributed to claims that open AI will have inherent beneficial effects on market competition⁷. However, access to market remains a constrained resource. Even well-resourced actors, who have the capital, talent and data to create large-scale models, do not always have an obvious way to deploy these models or ensure a return on investment, owing to substantial bottlenecks in market access, which–at present–runs through the large companies

by either cloud offerings or large-scale platform integrations. We see this in the example of 'open'Al company Mistral AI's decision to contract with Microsoft, allowing Microsoft to license a version of its Mixtral Large AI model to cloud customers through its Azure Cloud business. This is notable given that Mistral is one of the most well-financed AI start-ups building open models⁵⁷ and has marketed itself for its efficient use of computing⁵⁸. But even with these advantages, it still moved—alongside OpenAI and Inflection AI—to access the market through Microsoft's cloud platform⁵⁹.

3. Extensibility. Extensibility enables us to build on top of off-the-shelf models, fine-tuning them for one purpose or another. It is a key feature championed particularly by corporate actors invested in open AI⁶⁰. This is in large part because the work of 'extending' off-the-shelf models doubles as free product development for those who might want to repurpose a fine-tuned model. Extending an open AI model means that those doing this work do not start with a blank slate. They take a large model, already laboriously and expensively trained, and adjust its parameters and generally train it on further, often specialized, data in service of adapting its performance to a particular domain or task. Notable editorial decisions have already been made during the process of developing the 'base model⁶¹⁻⁶³.

The political economy of open AI

Here we review the materials—models, data, labour, frameworks, and computational power—frequently involved in creating and using large AI systems^{2,64}. This helps us evaluate which parts of these systems are or can be made open, which aren't or can't, and in what ways.

AI models

Much of the continuing discourse about open AI focuses on AI models, which are only one part of an operational AI system and which on their own do not account for the full development-to-deployment life cycle of an AI system. An AI model refers to an algorithmic system that has been trained and evaluated using large amounts of data to produce statistically likely outputs in response to a given input, stored as numerical weights. For example, ChatGPT works by applying generative pre-trained transformer (GPT) models, which were trained on huge amounts of text data, much of it scraped from the web. These GPT models are one part of ChatGPT's suite of client-specific software, which includes a web client and iOS and Android apps, each of which require discrete libraries and expertise to maintain, as well as skilled people to maintain them for as long as they exist⁴⁸. These clients incorporate GPT models as only one part of a user-facing interface. Once trained, an AI model can be released in the same way other software code would be released-under an open licence for reuse or otherwise made available online. Reusing an already-trained AI model does not require having access to the underlying training or evaluation data nor does it require weights or other system details be made available. In this sense, many AI systems that are labelled open are playing with the term loosely. Instead of providing meaningful documentation and access, they are effectively wrappers around closed models, inheriting undocumented data, failing to provide annotated reinforcement learning from human feedback (RLHF) training data and labour-process information and rarely publishing their findings, let alone documenting these in independently reviewed publications¹⁵.

There are now several examples of large-scale open AI models available for some degree of public reuse: these include Meta's LLaMA-2 (ref. 60) and LLaMA-3 (ref. 43); Falcon 40B, developed by the UAE's Technology Innovation Institute, trained on AWS⁶⁵; MosaicML's MPT⁶⁶ models and Mistral AI's Mixtral 8x22B, both tied to Microsoft's Azure; BigScience's BLOOM model, trained on the French Jean Zay supercomputer⁶⁷. Placing all of these under the singular label of open does a disservice to the serious distinctions between them and contributes to the confusion around the term. Companies such as Hugging Face and Stability AI offer open AI models to their customers and the public. Their business models rely not on licensing proprietary models themselves but instead on charging for extra features and labour on top of open models, features such as API access, model training on custom data and security and technical support as a paid service⁶⁸. 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 only on fostering research on large-scale AI⁶⁹, licensing its models under the very permissive Apache 2.0 open-source licence for use by AI researchers⁵⁶. Among those engaging in open AI, EleutherAI offers arguably the most maximally open AI systems.

A handful of academic projects have also produced large open AI models at smaller scales. These include Stanford's Alpaca model, well known for having been developed to run on a single laptop—a notable feat given the computationally intensive nature of deploying such models⁷⁰. However, even a chatbot based on this extremely computationally efficient model became too costly—and risky, owing to the model's 'hallucinations'—to continue running and the team has since taken it down⁷¹.

The present pattern in AI development takes a bigger-is-better approach when it comes to data, computing and model size³³. The bigger the model, the more resource intensive it is to train and calibrate and thus the more difficult it is to produce outside large technology companies. Although we know that the largest openly available AI model is at present LLaMA-3 and that it was trained on 15 trillion tokens⁴², information on the datasets for models has become increasingly opaque, for closed and ostensibly open models alike. OpenAI has not released the size of GPT-4 (ref. 72), Anthropic's technical report does not discuss the size of Claude 3's training data⁷³ and Mistral AI has declined to release even the size of the training data of its openly available model, citing the "highly competitive nature of the field"74. Further, although fine-tuning a model for a particular task or domain is less computationally expensive per instance (but much more environmentally costly in aggregate), such third parties can only build on top of models that they cannot scrutinize nor replicate, leading to an 'upper class of AI'³³.

Data

Data shaped to exacting (and labour-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 computing when building large-scale AI^{48,75}. Both are essential and, in the present 'rush-to-scale' pattern, the more of each, the 'better' these models perform^{33,76}.

Data are frequently a closed element of many AI offerings advertising themselves as open: many large-scale AI models described as open neglect to provide even basic information about the underlying data used to train the system⁷⁷, let alone offering the underlying training data openly or documenting its provenance. Lack of data transparency presents a serious challenge to any claims made around the benefits of open AI and hinders the kind of validation or reproducibility needed for sound science.

Scraping data to create datasets for AI development raises issues of extraction and intellectual property that are particularly relevant to concerns about concentration in the AI sector. 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, such as GhanaNLP⁷⁸ and Lesan AI⁷⁹. This means that, even though it is possible to train models without copyrighted content⁸⁰, 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 are being legally

contested⁸¹, and willing and able to weather the cost of lawsuits in either instance⁸². Legal or not, the practice of indiscriminately trawling web data to create systems that are now being poised to undercut the livelihoods of writers, artists and programmers—whose labour created such 'web' data in the first place—has raised alarm and ire⁸³, and lawsuits filed on behalf of these actors are now moving forward⁸⁴.

These concerns are particularly pressing considering the colonial echoes present in current data labour practices: AI systems frequently rely on data and labour resources from the majority world⁸⁵, and the founder of the GhanaNLP open-source project has noted that big tech's open source risks enabling continued colonial exploitation⁸⁶⁻⁸⁸. Such exploitation also runs directly counter to majority-world movements for data sovereignty, exemplified by projects such as Te Hiku Media, who point out 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"⁸⁹.

This is not an argument for closed datasets, which compound this issue. It is an entreaty to be clear about precisely what open datasets can and cannot 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 were specifically licensed for non-commercial use or were 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 licensed under the General Public License⁹⁰, an open-source licence that requires derivative code to be released under the same terms. However, even using permissively licensed code to train generative AI may similarly violate provisions requiring attribution, which current generative AI systems could, but do not, provide at present.

Datasets such as the Pile⁶¹ and Common Crawl^{92,93} are widely available but extra labour is required to make such datasets useful for the purpose of building large AI models. Careful curation and remixing of datasets is necessary to create performant AI: 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 noise⁶⁷. Although presumptively the larger datasets used by companies require proportionately similar levels of labour, we know little to nothing about them, even those that claim to be open.

Labour

The insatiable need of large-scale AI systems for curated, labelled, carefully organized data means that building AI at scale requires substantial human labour. This labour creates the 'intelligence' that AI systems are marketed as making computational^{94,95}. This labour can be roughly categorized as applied to:

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

Generative AI systems are trained and evaluated on a broad range of human-generated text, speech 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 to ensure that the outputs of the model stay within the bounds of 'acceptable'⁹⁶—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 from human feedback, or RLHF, which is a technical-sounding term that, in practice, refers to thousands of hours of human labour, during which 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, and their choices would be fed back into the system⁹⁷. Although data preparation and model calibration require extensive, rarely heralded labour that is fundamental in attaching meaning to the data that shape AI systems, companies generally release little if any information about the labour practices underpinning this data work, and failing to release such information is seldom criticized as a form of closedness. What we do know about these processes is largely the product of either investigative journalism⁹⁸⁻¹⁰⁰ or organizing by workers and researchers¹⁰¹⁻¹⁰³.

The labour required to curate, prepare data and calibrate systems is poorly paid but it still costs a large 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).

Development frameworks

Development frameworks make it easier for those developing software to build and deploy it in regimented, predictable and expedient ways. They are part of standard development practices and are not unique to AI. They work by providing pre-written pieces of code, templatized workflows, evaluation tools and other standardized methods 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. 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 deployed⁴.

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. There are many more pre-trained AI models that exclusively work within the PyTorch framework than there are those that work with TensorFlow^{33,104}. PyTorch is also the most popular framework in academic AI research, used in most academic papers^{105,106}.

PyTorch was initially developed for internal use by Meta but was released publicly in 2017. Although PyTorch operates as a research foundation under the umbrella of the Linux Foundation¹⁰⁷, it continues to be financially supported by Meta^{107,108} and its lead maintainers (responsible for governance and decision-making) are all Meta employees¹⁰⁹. TensorFlow was originally developed and released by Google Brain in 2015 (ref. 110) and continues to be directed and financially supported by Google, which also employs many of its core contributors¹¹¹.

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 companies developing them. Most notably, they allow Meta, Google and those steering framework development to standardize AI construction so that the results are compatible with their own company platforms—ensuring that their framework leads developers to create AI systems that, like Lego, snap into place with their own company systems¹¹². In the case of Meta, this allows them to more easily integrate and commercialize systems developed, tuned or deployed using PyTorch. Zuckerberg clearly stated these benefits to Meta in a 2023 earnings call, in which he said, "[PyTorch] has generally become the standard in the industry [...] 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."¹¹³ and reiterated this point in a 2024 earnings call³¹. This is true for Google and TensorFlow as well. In the case of Google, TensorFlow has been created to easily and intuitively operate with Google's Tensor Processing Unit (TPU) hardware, the powerful proprietary computing infrastructure core to Google's cloud AI computing business. This enables Google to optimize their commercial cloud offerings for AI development, positioning these products as the engine of AI. In this way, open development frameworks can work to entrench and bolster corporate AI dominance.

Open AI development frameworks can also allow those bankrolling and directing their development to create on-ramps to profitable computing and other service offerings. Similarly to how corporate representatives drive governance of internet standards to the exclusion of others¹¹⁴, AI companies shape the work practices of researchers and developers³³ such that new AI models can be easily integrated and commercialized. This gives the company offering the framework substantial 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^{4,115}.

Computational power

Developing large AI models requires massive datasets, which require massive computational power to process^{49,76}. Contemporary AI development is characterized by a race to scale³³, with older estimates showing that the amount of computing used to train models has increased about 300,000 times in 6 years, roughly an 8-fold increase each year¹¹⁶, and recent estimates of data use showing an increase in dataset size of around 2.4 times per year¹¹⁷. Access to computing remains a notable barrier to practical reusability for many open AI systems, because of the high cost involved in both training and running inference (51,686 kWh, 7,571 kWh and 1×10^{-4} kWh for training, fine-tuning and inference energy costs, respectively, in one case¹¹⁸) on large AI models at scale (that is, instrumenting them in a product or API for widespread public use). Furthermore, eking out maximal computational capacity from specialized hardware requires specialized and, in some cases, proprietary software systems.

It is hard to overstate Nvidia's dominance here: the company maintains a 70-90% market share for state-of-the-art AI chips¹¹⁹. Moreover, more than four million developers rely on CUDA¹²⁰, the 'de facto industry standard'¹²¹ partly proprietary framework developed by Nvidia that only supports training on the company's proprietary graphics processing units (GPUs) (specialized computer processors, originally developed for gaming, now primarily used for AI training because they allow many calculations to be performed quickly in parallel). The CUDA development ecosystem is a key element of Nvidia's powerful market dominance⁴⁹ (with the company's market share at 88% for GPUs¹²²) and has been nurtured and extended since 2006, giving it a big head start. Like Apple's developer ecosystem-which offers those wishing to build apps and services for the company's proprietary operating systems high-quality building blocks-CUDA provides expansive and norm-setting resources to AI researchers and developers.

In short, the computational resources needed to build new AI models and use existing ones at scale, outside privatized enterprise contexts and individual tinkering, are scarce, extremely expensive and concentrated among only a handful of corporations (with Nvidia at the helm^{49,122}), who themselves benefit from economies of scale, the capacity to control the software that optimizes computing and the ability sell costly access to computational resources^{33,123}. The seamlessness of integration across computational provision and model access is seen by some as powering demand for cloud infrastructure providers^{124,125}, further suggesting that it is ownership of an ecosystem, rather than the ability to produce a successful model or product offering, that determines competitiveness in AI.

Conclusion

By dissecting the pieces that together comprise modern AI systems and examining which of these pieces can and cannot be made open, we reveal a map of open AI which shows that, even at its most maximal, open AI is highly dependent on the resources of a few large corporate actors, who effectively control the AI industry and the research ecology beyond.

For this reason, the pursuit of even the most open AI will not on its own lead to a more diverse, accountable or democratized ecosystem. even though it may have other benefits. We also see that, as in the past, big tech companies vying for AI advantage are making use of open AI to consolidate market advantage while deploying the rhetorical wand of openness to deflect from accusations of AI monopoly and attendant regulation¹. The reality is, however open it is, when AI systems are deployed at scale across sensitive domains, they can have diffuse and profound effects that should not be determined by the small handful of for-profit companies who at present control the resources required to create and deploy these systems at scale, bringing them in front of the millions of customers who will be directly affected by them, particularly when these effects cannot be foreseen simply by examining system code, model weights and documentation. The creation of meaningful alternatives to the present AI model will not be accomplished through the pursuit of open AI development alone, even though elements such as data transparency and documentation are valuable for accountability, and maximally open AI projects helpfully illustrate the limits of what is possible. Focusing policy intervention on whether AI will be open or closed serves to distract from the overwhelmingly opaque nature of most corporate AI systems, both open and closed, in turn drawing valuable energy and initiative away from questions on the implications of AI in practice.

Unless pursued alongside other strong measures to address the concentration of power in AI, including antitrust enforcement and data privacy protections, the pursuit of openness on its own will be unlikely to yield much benefit. This is because the terms of transparency, and the infrastructures required for reuse and extension, will continue to be set by these same powerful companies, who will be unlikely to consent to meaningful checks that conflict with their profit and growth incentives.

We need a wider scope for AI development and greater diversity of methods, as well as support for technologies that more meaningfully attend to the needs of the public, not of commercial interests. And we need space to ask 'why AI' in the context of many pressing social and ecological challenges. Creating the conditions to make such alternatives possible is a project that can coexist with, and even be supported by, regulation. But pinning our hopes on 'open' AI in isolation will not lead us to that world, and—in many respects—could make things worse, as policymakers and the public put their hope and momentum behind open AI¹²⁶, assuming that it will deliver benefits that it cannot offer in the context of concentrated corporate power.

Data availability

We do not analyse nor supply a dataset, because our work does not rely on computational techniques.

- LeCun, Y. Testimony before the U.S. Senate Select Committee on Intelligence (2023).
 Luitse, D. & Denkena, W. The great transformer: examining the role of large language models
- in the political economy of AI. *Big Data Soc.* 8, 20539517211047734 (2021).
 Rikap, C. Same End By Different Means: Google, Amazon, Microsoft and Meta's Strategies
- Kikap, C. Same End By Different Means: Google, Amazon, Microsoft and Metas Strategies to Organize Their Frontier AI Innovation Systems. *City, University of London* https://www. cdh.cam.ac.uk/wp-content/uploads/2023/06/8.-Rikap-2023-Same-end-different-meanslonger-version-CITYPERC.pdf (2023).

- Whittaker, M. The steep cost of capture. *Interactions* 28, 50–55 (2021).
 Through a historical analysis of the military's influence over AI of the past, this piece shows how AI of the present depends on resources held by a few companies, thus allowing them to capture AI research agendas.
- Luitse, D. Platform power in AI: the evolution of cloud infrastructures in the political economy of artificial intelligence. *Internet Policy Rev.* 13, 1–44 (2024).
- O'Brien, M. White House wades into debate on 'open' versus 'closed' artificial intelligence systems. Associated Press https://apnews.com/article/ai-executive-order-biden-opensourcemodels-1c42092e55729d731d246440094f7fed (2024).
- Staff in the Office of Technology. On Open-Weights Foundation Models. Federal Trade Commission https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/07/openweights-foundation-models (2024).
- Competition & Markets Authority. CMA AI strategic update. UK Government https://www. gov.uk/government/publications/cma-ai-strategic-update/cma-ai-strategic-update (2024).
- Bommasani, R. et al. Considerations for Governing Open Foundation Models. Stanford University https://hai.stanford.edu/issue-brief-considerations-governing-open-foundationmodels (2023).
- Karabus, J. GitHub CEO says EU AI Act shouldn't apply to open source devs. The Register https://www.theregister.com/2023/02/07/github_ceo_ai_act/ (2023).
- Deutsch, J. OpenAI Backs Idea of Requiring Licenses for Advanced AI Systems. Bloomberg https://www.bloomberg.com/news/articles/2023-07-20/internal-policy-memo-showshow-openai-is-willing-to-be-regulated (2023).
- Kapoor, S. et al. Position: On the societal impact of open foundation models. Proc. Machine Learning Res. 235, 23082–23104 (2024).
- Seger, E. et al. Open-sourcing highly capable foundation models: an evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. Preprint at http:// arxiv.org/abs/2311.09227 (2023).
- Bommasani, R. et al. The Foundation Model Transparency Index. Preprint at https://doi. org/10.48550/arXiv.2310.12941 (2023).
- 15. Liesenfeld, A., Lopez, A. & Dingemanse, M. in Proc. 5th International Conference on
- Conversational User Interfaces 1–6 https://doi.org/10.1145/3571884.3604316 (ACM, 2023).
 Suchman, L. The uncontroversial 'thingness' of Al. Big Data Soc. 10, 20539517231206794 (2023).

An important short piece problematizing the stability of the notion of 'AI'.

- Gent, E. The tech industry can't agree on what open-source AI means. That's a problem. MIT Technology Review https://www.technologyreview.com/2024/03/25/1090111/techindustry-open-source-ai-definition-problem/ (2024).
- Open Source Initiative. The Open Source AI Definition 1.0. Open Source Initiative https:// opensource.org/ai/open-source-ai-definition (2024).
- Coleman, E. G. Coding Freedom: The Ethics and Aesthetics of Hacking (Princeton Univ. Press, 2013).

An ethnography examining the politics and ethics of free and open source software communities.

- 20. Kelty, C. M. Two Bits: The Cultural Significance of Free Software (Duke Univ. Press, 2008).
- Meneely, A. & Williams, L. in Proc. 16th ACM Conference on Computer and Communications Security 453–462 https://doi.org/10.1145/1653662.1653717 (ACM, 2009).
- Rosenberg, J. The meaning of open. Google Public Policy Blog https://publicpolicy. googleblog.com/2009/12/meaning-of-open.html (2009).
- 23. Alexy, O. Free Revealing: How Firms Can Profit From Being Open (Springer, 2009).
- IBM. A strong history and commitment to open source. IBM https://www.ibm.com/ opensource/story/.
- AT.40099 Google Android. European Commission https://competition-cases.ec.europa.eu/ cases/AT.40099 (2015).
- Barr, J. Amazon DocumentDB (with MongoDB Compatibility): Fast, Scalable, and Highly Available. AWS News Blog https://aws.amazon.com/blogs/aws/new-amazon-documentdbwith-mongodb-compatibility-fast-scalable-and-highly-available/ (2019).
- Staff in the Office of Technology. Cloud Computing: Taking Stock and Looking Ahead. Federal Trade Commission https://www.ftc.gov/news-events/events/2023/05/cloudcomputing-taking-stock-looking-ahead (2023).
- McLaughlin, K. & Gardizy, A. After Years of Resistance, AWS Opens Checkbook for Open-Source Providers. *The Information* https://www.theinformation.com/articles/afteryears-of-resistance-aws-opens-checkbook-for-open-source-providers (2023).
- Asay, M. & Mehla, A. MongoDB and AWS Expand Global Collaboration. MongoDB https://www.mongodb.com/blog/post/mongodb-aws-expand-global-collaboration (2022).
- Zuckerberg, M. Open Source AI Is the Path Forward. Meta https://about.fb.com/news/ 2024/07/open-source-ai-is-the-path-forward/ (2024).
- 31. Zuckerberg, M. & Dorell, K. Fourth Quarter 2023 Results conference call (2024).
- Ojewale, V., Steed, R., Vecchione, B., Birhane, A. & Raji, I. D. Towards AI accountability infrastructure: gaps and opportunities in AI audit tooling. Preprint at https://doi.org/ 10.48550/arXiv.2402.17861 (2024).
- 33. Gururaja, S., Bertsch, A., Na, C., Widder, D. & Strubell, E. in Proc. 2023 Conference on Empirical Methods in Natural Language Processing (eds Bouamor, H., Pino, J. & Bali, K.) 13310–13325 https://doi.org/10.18653/v1/2023.emnlp-main.822 (Association for Computational Linguistics, 2023). This is an interview study of academic NLP researchers on the incentives governing

their field and on the effect of corporate power in their practice.

- 34. Khan, L. M. Sources of tech platform power. Georgetown Law Technol. Rev. 2, 325 (2018).
- Vipra, J. & Korinek, A. Market concentration implications of foundation models: the invisible hand of ChatGPT. Brookings Center on Regulation and Markets https://www.brookings.edu/ articles/market-concentration-implications-of-foundation-models-the-invisible-hand-ofchatgpt/ (2023).
- Ezrachi, A. & Stucke, M. E. Virtual Competition: The Promise and Perils of the Algorithm-Driven Economy (Harvard Univ. Press, 2019).
- OECD. Artificial intelligence, data and competition. OECD https://www.oecd.org/en/ publications/artificial-intelligence-data-and-competition_e7e88884-en.html (2024).

- Competition and Markets Authority, UK Government. AI Foundation Models Technical Update Report. UK Government https://www.gov.uk/cma-cases/ai-foundation-modelsinitial-review (2024).
- West, S. M. & Kak, A. Al Now 2023 Landscape: Confronting Tech Power. Al Now Institute https://ainowinstitute.org/2023-landscape (2023).
- Nathan, A., Grimberg, J. & Rhodes, A. Gen Al: Too Much Spend, Too Little Benefit? Goldman Sachs https://www.goldmansachs.com/intelligence/pages/gs-research/ gen-ai-too-much-spend-too-little-benefit/report.pdf (2024).
- Solaiman, I. in Proc. 2023 ACM Conference on Fairness, Accountability, and Transparency 111–122 https://doi.org/10.1145/3593013.3593981 (ACM, 2023).
- This is a key piece that was the first to argue that Al is not 'binary' in an Al context.
 42. Meta. Introducing Meta Llama 3: The most capable openly available LLM to date. Meta Al https://doi.org/10.1016/j.00011
- https://ai.meta.com/blog/meta-llama-3/ (2024). 43. Meta Meta Llama 3 License. *Meta* https://llama.meta.com/llama3/license/ (2024).
- Meta Meta Lama S Elcense: Meta https://tama.neta.com/tamao/tecise/ (2024).
 Maffulli, S. Meta's LLaMa 2 license is not Open Source. Open Source Initiative https://
- opensource.org/blog/metas-llama-2-license-is-not-open-source (2023). 45. Nolan, M. Llama and ChatGPT Are Not Open-Source. *IEEE Spectrum* https://spectrum.
- ieee.org/open-source-llm-not-open (2023).
 Hull, C. Is Llama 2 open source? No and perhaps we need a new definition of open... OpenSource Connections https://opensourceconnections.com/blog/2023/07/19/ is-llama-2-open-source-no-and-oerhaps-we-need-a-new-definition-of-open/(2023).
- Buchanan, B. The AI Triad and What It Means for National Security Strategy. Center for Security and Emerging Technology https://cset.georgetown.edu/publication/the-ai-triadand-what-it-means-for-national-security-strategy/ (2020).
- Musser, M., Gelles, R., Kinoshita, R., Aiken, C. & Lohn, A. The Main Resource Is the Human. Center for Security and Emerging Technology https://cset.georgetown.edu/publication/ the-main-resource-is-the-human/ (2023).
- Vipra, J. & West, S. M. Computational Power and Al. Al Now Institute https://ainowinstitute. org/publication/policy/compute-and-ai (2023).
- Dettmers, T., Pagnoni, A., Holtzman, A. & Zettlemoyer, L. QLoRA: Efficient Finetuning of Quantized LLMs. Proc. 37th Int. Conf. Neural Inf. Process. (NIPS '23) Article 441, 10088–10115 (Curran Assoc., Red Hook, NY, 2024).
- Birhane, A., Prabhu, V., Han, S., Boddeti, V. N. & Luccioni, A. S. Into the LAION's Den: Investigating Hate in Multimodal Datasets. 37th Conf. Neural Inf. Process. Syst. Datasets Benchmarks Track https://openreview.net/forum?id=6URyQ9QhYv (2023).
- 52. Gerchick, M. et al. in Proc. 2023 ACM Conference on Fairness, Accountability, and Transparency 1292–1310 https://doi.org/10.1145/3593013.3594081 (ACM, 2023).
- 53. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **1**, 206–215 (2019).
- Lipton, Z. C. The mythos of model interpretability: in machine learning, the concept of interpretability is both important and slippery. *Queue* 16, 31–57 (2018).
- Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Wortman Vaughan, J. W. & Wallach, H. in Proc. 2021 CHI Conference on Human Factors in Computing Systems 1–52 https://doi.org/10.1145/3411764.3445315 (ACM, 2021).
- EleutherAI. License EleutherAI/gpt-neox. GitHub https://github.com/EleutherAI/ gpt-neox?tab=Apache-2.0-1-ov-file (2024).
- Metz, C. Mistral, French A.I. Start-Up, Is Valued at \$2 Billion in Funding Round. *The New York Times* https://www.nytimes.com/2023/12/10/technology/mistral-ai-funding.html (2023).
- 58. Mistral AI. https://mistral.ai/.
- Boyd, E. Introducing Mistral-Large on Azure in partnership with Mistral AI. Microsoft Azure Blog https://azure.microsoft.com/en-us/blog/microsoft-and-mistral-ai-announce-newpartnership-to-accelerate-ai-innovation-and-introduce-mistral-large-first-on-azure/ (2024).
- Touvron, H. et al. Llama 2: open foundation and fine-tuned chat models. Preprint at https://doi.org/10.48550/arXiv.2307.09288 (2023).
- Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. in Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency 610–623 https://doi.org/10.1145/ 3442188.3445922 (ACM, 2021).

This article documents that the creation of massive training datasets renders them difficult to scrutinize and increasingly opaque.

- Bertsch, A. et al. in Proc. 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP) (eds Hardmeier, C., Basta, C., Costa-jussà, M. R., Stanovsky, G. & Gonen, H.) 235–243 https://doi.org/10.18653/v1/2022.gebnlp-1.24 (Association for Computational Linguistics, 2022).
- Longpre, S. et al. A Pretrainer's Guide to Training Data: Measuring the Effects of Data Age, Domain Coverage, Quality, & Toxicity. Proc. 24 Conf. N. Am. Ch. Assoc. Comp. Linguistics Hum. Language Tech. Vol. 1, 3245–3276 (Assoc. Comp. Linguistics, Mexico City, Mexico, 2023).
- Widder, D. G. & Nafus, D. Dislocated accountabilities in the "AI supply chain": modularity and developers' notions of responsibility. *Big Data* Soc. 10, 20539517231177620 (2023).
- Almazrouei, E. et al. The Falcon series of open language models. Preprint at https://doi. org/10.48550/arXiv.2311.16867 (2023).
- The MosaicML NLP Team. Introducing MPT-7B: A New Standard for Open-Source, Commercially Usable LLMs. *Databricks* https://www.databricks.com/blog/mpt-7b (2023).
 Scao, T. L. et al. BLOOM: a 176B-parameter open-access multilingual language model.
- Ostado, F. E. et al. B. E. Orazin, a Frob parameter oper faces and thinking at tanguage model.
 Preprint at https://arxiv.org/abs/2211.05100 (2023).
 Osman, L. & Sewell, D. Hugging Face. Contrary Research https://research.contrary.com/
- Osman, L. & Sewell, D. Hugging Face. Contrary Research https://research.contrary.com/ company/hugging-face (2022).
- 69. EleutherAI. EleutherAI About. EleutherAI https://www.eleuther.ai/about.
- Edwards, B. You can now run a GPT-3-level AI model on your laptop, phone, and Raspberry Pi. Ars Technica https://arstechnica.com/information-technology/2023/03/ you-can-now-run-a-gpt-3-level-ai-model-on-your-laptop-phone-and-raspberry-pi/ (2023).
- Germain, T. Stanford Researchers Take Down Alpaca AI Due to 'Hallucinations' and Rising Costs. *Gizmodo Australia* https://gizmodo.com/stanford-ai-alpaca-llama-facebook-takendown-chatgpt-1850247570 (2023).
- 72. OpenAI. GPT-4 technical report. Preprint at http://arxiv.org/abs/2303.08774 (2023).

- Anthropic. The Claude 3 Model Family: Opus, Sonnet, Haiku. Anthropic https:// www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_ Card_Claude_3.pdf.
- Mensch, A. mistralai/Mistral-7B-v0.1. Training data? Hugging Face https://huggingface.co/ mistralai/Mistral-7B-v0.1/discussions/8 (2023).
- 75. Xu, H. et al. Demystifying CLIP data. Preprint at https://arxiv.org/abs/2309.16671 (2023).
- Kaplan, J. et al. Scaling laws for neural language models. Preprint at https://doi.org/ 10.48550/arXiv.2001.08361 (2020).
- Gebru, T. et al. Datasheets for datasets. Commun. ACM 64, 86–92 (2021).
 GhanaNLP. About Us. GhanaNLP https://ghananlp.org/about.
- GhanaNLP. About Us. GhanaNLP https://ghananlp.org/about.
 Deck, A. The AI startup outperforming Google Translate in Ethiopian languages. Rest of
- World https://restofworld.org/2023/3-minutes-with-asmelash-teka-hadgu/ (2023).
 Knibbs, K. Here's Proof You Can Train an AI Model Without Slurping Copyrighted Content. WIRED https://www.wired.com/story/proof-you-can-train-ai-without-slurping-copyrighted-
- content/ (2024).
 81. Khan, M. & Hanna, A. The subjects and stages of AI dataset development: a framework for dataset accountability. *Ohio State Law J.* **19**, 171 (2023).
- Smith, B. & Nowbar, H. Microsoft announces new Copilot Copyright Commitment for customers. *Microsoft On the Issues* https://blogs.microsoft.com/on-the-issues/2023/ 09/07/copilot-copyright-commitment-ai-legal-concerns/ (2023).
- Jiang, H. H. et al. in Proc. 2023 AAAI/ACM Conference on AI, Ethics, and Society 363–374 https://doi.org/10.1145/3600211.3604681 (ACM, 2023).
- Setty, R. Sarah Silverman, Authors Hit OpenAl, Meta With Copyright Suits. Bloomberg Law https://news.bloomberglaw.com/ip-law/sarah-silverman-authors-hit-openai-meta-withcopyright-suits (2023).
- Browne, G. Al Is Steeped in Big Tech's 'Digital Colonialism'. WIRED https://www.wired. com/story/abeba-birhane-ai-datasets/ (2023).
- Azunre, P. [@pazunre]. 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]. Twitter https://twitter.com/pazunre/status/1569743778524680192 (2022).
- Abebe, R. et al. in Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency 329–341 https://doi.org/10.1145/3442188.3445897 (ACM, 2021).
 Común, T. Resisting Data Colonialism – A Practical Intervention (Institute of Network).
- Común, T. Resisting Data Colonialism A Practical Intervention (Institute of Network Cultures, 2023).
- Te Hiku Media. Kaitiakitanga-License/LICENSE.md. GitHub https://github.com/TeHikuMedia/ Kaitiakitanga-License/blob/tumu/LICENSE.md (2018).
- Vincent, J. The lawsuit that could rewrite the rules of AI copyright. The Verge https://www. theverge.com/2022/11/8/23446821/microsoft-openai-github-copilot-class-action-lawsuit-aicopyright-violation-training-data (2022).
- Gao, L. et al. The Pile: an 800GB dataset of diverse text for language modeling. Preprint at https://doi.org/10.48550/arXiv.2101.00027 (2020).
- Baack, S. A in Proc. 2024 ACM Conference on Fairness, Accountability, and Transparency 2199–2208 https://doi.org/10.1145/3630106.3659033 (ACM, 2024).
- Luccioni, A. S. & Viviano, J. D. What's In The Box? A Preliminary Analysis of Undesirable Content in the Common Crawl Corpus. Proc. 59th Annu. Meeting Assoc. Comp. Linguistics and 11th Int. Joint Conf. Natural Lang. Processing Vol. 2 182–189 (Assoc. Comp. Linguistics, 2021).
- Taylor, A. The Automation Charade. Logic(s) Magazine https://logicmag.io/failure/theautomation-charade/ (2018).
- Williams, A., Miceli, M. & Gebru, T. The Exploited Labor Behind Artificial Intelligence. Noema Magazine https://www.noemamag.com/the-exploited-labor-behind-artificialintelligence/ (2022).
- Thylstrup, N. & Talat, Z. Detecting 'Dirt' and 'Toxicity': Rethinking Content Moderation as Pollution Behaviour. SSRN https://doi.org/10.2139/ssrn.3709719 (2020).
- Christiano, P. et al. Deep Reinforcement Learning From Human Preferences. Proc. 31 Int. Conf. Neural Inf. Processing Syst. (NIPS '17) 4302–4310 (Curran Assoc., Red Hook, NY, 2017).
- Perrigo, B. Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic. *Time.com* https://time.com/6247678/openai-chatgpt-kenya-workers/ (2023).
- Hao, K. The Hidden Workforce That Helped Filter Violence and Abuse Out of ChatGPT. The Wall Street Journal https://www.wsj.com/podcasts/the-journal/the-hidden-workforcethat-helped-filter-violence-and-abuse-out-of-chatgpt/ffc2427f-bdd8-47b7-9a4b-27e7267cf413 (2023).
- Hao, K. & Hernández, A. P. How the Al industry profits from catastrophe. MIT Technology Review https://www.technologyreview.com/2022/04/20/1050392/ai-industry-appenscale-data-labels/ (2022).
- Perrigo, B. 150 African Workers for ChatGPT, TikTok and Facebook Vote to Unionize at Landmark Nairobi Meeting. *Time.com* https://time.com/6275995/chatgpt-facebookafrican-workers-union/ (2023).
- 102. Mutemi, M. [@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 @OpenAl and @samasource got away with such exploitation and to urgently regulate tech work. https://t.co/9seeyGKqFM. Twitter https://twitter.com/MercyMutemi/status/1678984336996028416 (2023).
- Irani, L. C. & Silberman, M. S. in Proc. SIGCHI Conference on Human Factors in Computing Systems 611–620 https://doi.org/10.1145/2470654.2470742 (ACM, 2013).
- Foster, K. PyTorch vs TensorFlow: Who has More Pre-trained Deep Learning Models? HackerNoon https://hackernoon.com/pytorch-vs-tensorflow-who-has-more-pre-traineddeep-learning-models (2022).
- He, H. The State of Machine Learning Frameworks in 2019. The Gradient https://thegradient. pub/state-of-ml-frameworks-2019-pytorch-dominates-research-tensorflow-dominates-in dustry/ (2019).
- Paszke, A. et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Proc. 33 Int. Conf. Neural Inf. Processing Syst. Article 721, 8026–8037 (Curran Assoc., Red Hook, NY, 2019).

- 107. Meta. Announcing the PyTorch Foundation: A new era for the cutting-edge AI framework. Meta https://ai.meta.com/blog/pytorch-foundation/ (2022).
- Tarantola, A. Meta is spinning off the Pytorch framework into its own AI research foundation. *Engadget* https://www.engadget.com/meta-is-spinning-off-its-pytorch-ai-framework-into-its-own-research-foundation-140051987.html (2022).
- PyTorch Contributors. PyTorch Governance: Maintainers. PyTorch 2.2 documentation. PyTorch https://pytorch.org/docs/stable/community/persons_of_interest.html (2023).
- Metz, C. Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine. WIRED https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/ (2015).
- Contributors to tensorflow/tensorflow. GitHub https://github.com/tensorflow/tensorflow/ graphs/contributors (2024).
- Engler, A. How open-source software shapes AI policy. Brookings https://www.brookings. edu/articles/how-open-source-software-shapes-ai-policy/ (2021).
- Zuckerberg, M. & Crawford, D. First Quarter 2023 Results conference call (2023).
 Cath-Speth, C. J. N. Changing Minds and Machines: A Case Study of Human Rights Advocacy
- in the Internet Engineering Task Force (IETF). Thesis, Univ. Oxford (2021).
- Langenkamp, M. & Yue, D. N. in Proc. 2022 AAAI/ACM Conference on AI, Ethics, and Society 385–395 https://doi.org/10.1145/3514094.3534167 (ACM, 2022).

This study examines the economic value of open-source AI development. 116. Schwartz, R., Dodge, J. & Smith, N. A. Green AI. Commun. ACM 63, 54–63 (2020).

- Villalobos, P. et al. Will we run out of data? Limits of LLM scaling based on humangenerated data. Preprint at http://arxiv.org/abs/2211.04325 (2024).
- Luccioni, S., Jernite, Y. & Strubell, E. in Proc. 2024 ACM Conference on Fairness, Accountability, and Transparency 85–99 https://doi.org/10.1145/3630106.3658542 (ACM, 2024).

This paper measures the carbon cost (and therefore, indirectly, the US\$ cost) of different AI models, thus giving a scale of the resources they require to run.

- Leswing, K. Nvidia dominates the AI chip market, but there's more competition than ever. CNBC https://www.cnbc.com/2024/06/02/nvidia-dominates-the-ai-chip-market-but-theresrising-competition-.html (2024).
- Cherney, M. A. Exclusive: Behind the plot to break Nvidia's grip on AI by targeting software. Reuters https://www.reuters.com/technology/behind-plot-break-nvidias-grip-ai-bytargeting-software-2024-03-25/ (2024).
- Why do Nvidia's chips dominate the AI market? The Economist https://www.economist. com/the-economist-explains/2024/02/27/why-do-nvidias-chips-dominate-the-ai-market (2024).
- 122. Dow, R. Shipments of graphics add-in boards decline in Q1 of 24 as the market experiences a return to seasonality. Jon Peddie Research https://www.jonpeddie.com/news/shipmentsof-graphics-add-in-boards-decline-in-q1-of-24-as-the-market-experiences-a-return-toseasonality/ (2024).
- Patel, D. & Nishball, D. Google Gemini Eats The World Gemini Smashes GPT-4 By 5X, The GPU-Poors. Semianalysis https://www.semianalysis.com/p/google-gemini-eats-the-worldgemini (2023).
- Dotan, T. Microsoft Earnings Growth Accelerates on Stronger-Than-Expected Cloud Demand. The Wall Street Journal https://www.wsj.com/tech/microsoft-msft-q1-earningsreport-2024-b19e51eb (2023).
- Vyas, I. Al and Machine Learning Integration in SaaS Applications. *IEEE Computer Society* https://www.computer.org/publications/tech-news/trends/ai-and-machine-learningintegration/ (2023).
- Lowe, L. AI, Startups, & Competition: Shaping California's Tech Future. Chaindesk https:// www.chaindesk.ai/tools/youtube-summarizer/ai-startups-and-competition-shapingcalifornia-s-tech-future-pieVtTrbDBs (2024).

Acknowledgements Our analysis builds on the work of past scholars who have written on connections between open-source software, AI and corporate power: N. Dyer-Witheford, A. M. Kjosen, J. Steinhoff, S. Gürses, J. Cobbe, J. Singh, M. Veale, G. Coleman, C. Kelty, D. Nafus, I. Solaiman and P. Chavez. We are grateful to P. Azunre, A. Bertsch, S. Biderman, F. Chollet, S. Gururaja, A. Hanna, A. Kak, C. Kelty, J. Lund, K. Mahelona, V. Mathur, M. Mitchell, A. Skowron, J. Sundwall, Z. Talat and L. Villa for feedback and conversations that deeply enriched this paper. D.G.W. acknowledges the Digital Life Initiative.

Author contributions D.G.W. analysed the history of open source and divergent corporate stances, contributed data collection and analysis of AI resource dependencies, as well as writing and editing. M.W. contributed the taxonomy of AI systems that enabled an analysis of AI resource dependencies, historical data and analysis of open source, the three affordances of open AI, as well as writing and editing. S.M.W. contributed the market and competition analysis, outlined the material dimensions of open AI and implications for AI policy, as well as writing, editing and data collection.

Competing interests The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to David Gray Widder. Peer review information Nature thanks Alexander Brem, Jesus Gonzalez Barahona and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Reprints and permissions information is available at http://www.nature.com/reprints. Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author selfarchiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature Limited 2024