"One of Silicon Valley's Most Divisive Topics": How the Media Discusses Openness in AI

Tamara Paris, Jin L.C. Guo*, AJung Moon*

McGill University, Montréal, QC, Canada
tamara.paris@mail.mcgill.ca, jguo@cs.mcgill.ca, ajung.moon@mcgill.ca

Abstract

In the last few years, there has been a rapid release of highly capable Artificial Intelligence (AI) systems labeled as "open" or "open source." Openness in AI encompasses a wider range of practices than open source software: from AI systems with all components published publicly to those that only share model weights and prohibit certain uses. Today, embracing this concept is a contentious topic. Some argue that the benefits of openness outweigh its risks. Others believe open release of AI poses more risks to society. This "open vs. closed AI" debate has made headlines in major news outlets and has become a political topic. In this study, we analyze how AI openness is framed in the media through a qualitative analysis of 223 newspaper articles from the U.S., France, and China. We find that effective communication around this debate is hindered by inaccurate use of terminology, the presence of misleading information, and an either-or dichotomy of AI openness that does not reflect the complexity of AI system development. Our analysis also reveals significant heterogeneity across news sources, and that media discourse on AI openness focuses solely on a handful of models. The media can influence public opinion on AI, promote user adoption of certain models, and subsequently impact technology policy decisions. Therefore, we call on the AI community to help add clarity to the debate.

1 Introduction

The "open vs. closed AI" debate has received much attention in recent years. This debate refers to the opposition between those who embrace making AI available as openly as possible and those who find unrestricted openness practice too risky for AI. Openness in AI is frequently referred to as the accessibility of the AI system's components, and the ability for anyone to freely use, modify, study and share these components (White et al. 2024; Open Source Initiative 2024). This definition is adapted from open source software; however, due to the complexities of modern day AI architecture, such a definition does not directly translate to openness in AI (Paris, Moon, and Guo 2025; Choksi, Mandel, and Benthall 2025). As more capable AI models started to become publicly available in the last five years, scholars have identified benefits and assessed risks of openness in AI beyond those already discussed for software (Seger et al. 2023; Whittlestone and Ovadya 2020; Shevlane and Dafoe 2020). Their work encourages more thoughtful approach for model release (Chan et al. 2023) and alerts developers of potential misuses of their innovation (Anderljung, Hazell, and von Knebel 2025).

Despite the risks, however, openness in AI can offer regulatory advantages for the AI industry. For example, AI systems released under a free or open source license are currently excluded from the regulatory requirements of the EU AI Act (European Union 2024). General purpose AI systems released openly have fewer documentation requirements in the Act as well (Liesenfeld and Dingemanse 2024). Moreover, some governments and organizations advocate for openness in AI publicly: the United Nations AI Advisory Body, for example, underscores open sharing of AI models as an important factor to increasing global AI capacity (Nationen 2024). However, since 2023, some industry players have become more cautious about openness in AI, and choose to close their model for ethics, security, and profit (Claude AI Hub 2024) while others actively support open practices (Zuckerberg 2024; Clegg 2024).

To date, a number of major news outlets have published dedicated articles on this debate (Isaac 2024; Lin 2024b; Mims 2025), and several industry leaders have publicly taken a stance through sponsored articles and opinion pieces (Meta 2024; Lee 2023). The frame used by the media to describe AI can impact the public's perception of AI (Choi 2024; Brewer et al. 2022; Cui and Wu 2021). News coverage on the "open vs. closed AI" debate is no exception. Public perception of AI corporations can be shaped by how the debate is portrayed in the media; this can, in turn, affect both public adoption of certain models over others as well as AI policy decisions made by different institutions.

Therefore, in this study, we explore media coverage of AI openness in the United States, France, and China—three host countries of large AI models claiming to be open source. We qualitatively analyzed 223 newspaper articles from *The Wall Street Journal (WSJ)*, *The New York Times (NYT)*, *The Washington Post (WP)*, *USA Today*, *Le Monde*, and *China Daily*; we examined how the articles define openness in AI and what angles they adopt in discussing the associated risks and benefits. Our analysis highlights the misuse of terminologies (e.g., "open source software") and polar-

*These authors contributed equally.

ized stances of the debate the current news tends to take. Our result reflects the growing geopolitical and economic significance of AI openness and underscores the urgent need for the broader AI community to better engage the public and journalists in the debate.

2 Background

Our work is situated within two broader contexts: (1) the evolution of openness in AI across industry and research (Sec. 2.1) and benefits and risks identified by the scholars (Sec. 2.2); (2) AI in the media (Sec. 2.3).

2.1 The Evolution of Openness in AI

Building on decades of history of open source in the software industry, open source became common practice in the AI community. In the 2010s, many models started to be released as open source alongside research articles, including those developed by large companies (e.g. BERT (Devlin et al. 2019), ResNet (He et al. 2016)). Numerous deep learning frameworks were also released openly, such as TensorFlow in 2015 (Abadi et al. 2015), followed by PyTorch, and CNTK (Langenkamp 2022).

As models became increasingly capable and large around 2020, there was a wider diversity in the way AI models were released (Solaiman 2023; Liesenfeld, Lopez, and Dingemanse 2023; Liesenfeld and Dingemanse 2024). For instance, some of the most popular models reveal little information about their training process and do not provide access to their weights, such as Google's PaLM and Bard, OpenAI's GPT-3, GPT-4, and ChatGPT, and Anthropic's Claude (Gao et al. 2023). Other models such as Meta's Llama 2 and Mistral AI's Mistral 7B have their weights released publicly but not their training code or data (Liesenfeld and Dingemanse 2024). Yet others release their dataset, codes and model parameters publicly such as OLMo (Groeneveld et al. 2024). Even among these fully accessible models, some are released with licenses that restrict certain downstream uses for responsive motives, such as the OpenRAIL licenses (Muñoz Ferrandis and Contractor 2023). BLOOM is an example of one such model (Workshop et al. 2023). With the absence of a shared definition and expectation of openness in AI, many companies have been accused of participating in openwashing practices where they label their models as "open" despite it having restrictive licenses or lacking transparency around training data (Kessler 2024; Liesenfeld and Dingemanse 2024; Maffulli 2023).

The wide range of model release methods and the growing discourse around the ethics of openly released models underscored the need for the AI community to critically examine what openness should mean for AI. Since 2022, several initiatives aimed to define AI openness. While some choose to take a binary approach—considering an AI to be either open or not (Open Source Initiative 2024) analogous to the binary definition of open source software (Perens et al. 1999)—others define it as a spectral and composite notion (Liesenfeld and Dingemanse 2024; Nobel, Rozenshtein, and Sharma 2025; Solaiman 2023; White et al. 2024; Choksi, Mandel, and Benthall 2025). The latter approach aims to encompass the multitude of access styles (Bucknall and Trager 2023; Kembery, Bucknall, and Simpson 2024) and the possible mix of open and closed-ness of the components comprising an AI system (Basdevant et al. 2024). Ding et al. advocate for embracing openness practices beyond access to AI components and call for more open collaboration (Ding et al. 2023). Choksi, Mandel, and Benthall Others draw inspiration from other disciplines to broaden the scope of openness in AI (Paris, Moon, and Guo 2025). To date, the boundaries of openness in AI still remain underdefined, creating a void for interested actors to fill in ways that benefit them. Given the lack of agreed-upon definition of open source AI, we use *openness in AI* in this article to refer to models that are framed as "open."

2.2 Risks and Benefits of Openness Practices

Many scholars documented the benefits and risks of AI openness, both for the open release of models (Seger et al. 2023; Eiras et al. 2024) and for the open access to AI research (Whittlestone and Ovadya 2020; Shevlane and Dafoe 2020). Some argue that the AI community and policy makers should generally support AI openness (Nobel, Rozenshtein, and Sharma 2025; Shrestha, von Krogh, and Feuerriegel 2023; Eiras et al. 2024), and that research should prioritize the use of open models (Palmer, Smith, and Spirling 2024). Others adopt a more cautious approach and propose alternatives to openness and advocate for mitigation measures (Seger et al. 2023; Chan et al. 2023).

Here, we present a brief summary of the benefits and risks of openness practices in AI discussed in previous work.

Societal Benefits. Previous studies mostly focus on the following five societal benefits of AI openness: (1) **advancing research, progress and innovation**, notably by offering greater extensibility, customizability, and interoperability (Shrestha, von Krogh, and Feuerriegel 2023; Eiras et al. 2024); (2) **fostering AI adoption**, notably by reducing cost of AI technology for end-users (Eiras et al. 2024); (3) **democratizing AI development and redistributing influence over downstream integration** by allowing each community to adapt AI models to their needs, empowering developers (Seger et al. 2023; Eiras et al. 2024), and reducing barriers to entry for a broader range of individuals (Nobel, Rozenshtein, and Sharma 2025); (4) **enhancing transparency and accountability** by enabling public audit (Nobel, Rozenshtein, and Sharma 2025), thereby increasing public trust and reducing copyright disputes (Eiras et al. 2024); (5) **helping to identify bugs and safety issues** through independent evaluations (Seger et al. 2023), allowing external validation of system integrity and continuous improvement of AI safeguards (Nobel, Rozenshtein, and Sharma 2025).

However, some of these benefits are direct translations of benefits from open source practices in the software domain and are not necessarily grounded in empirical evidence. For example, the argument that openness in AI can improve safety by enabling independent audits may not hold true in practice. Auditing large scale models can require resources beyond the reach of any independent auditor regardless of

whether they have access to AI components (Widder, West, and Whittaker 2023). Openness might also fail to meaningfully redistribute AI power; instead, it could reinforce existing power dynamics by locking in developers to a monopoly of an AI model and its ecosystem (Seger et al. 2023; Widder, West, and Whittaker 2023), particularly without careful structuring of AI governance (Nobel, Rozenshtein, and Sharma 2025).

Other benefits depend on the definition of openness in AI. For example, the ability to resolve copyright disputes mainly relies on whether the training data is open: even though there are efforts to demonstrate a model's use of copyrighted data, such as through membership inference attacks, these may not constitute sufficient legal evidence (Zhang et al. 2025).

Corporate Benefits. There are many reasons why companies would make their models publicly accessible. Osborne shows that open source practices can help the company **build a good reputation**, **established their dominance** in the market by shaping industry standards and increasing adoption of their model, and **improve the quality of its model** (Osborne 2024). Others suggest that it helps **reduce costs and accelerate development** by benefiting from external contributors and recruiting developers who are already familiar with system (Widder, West, and Whittaker 2023; Seger et al. 2023).

Societal Risks. One of the main concerns about open release of large AI models involves the increased **risk of misuse** of publicly accessible models. This includes the use of these models to create scams, cyber attacks, bio-weapons, disinformation, non-consensual sexual imagery, to name a few (Widder et al. 2022; Seger et al. 2023; Kapoor et al. 2024; Chan et al. 2023). The fundamental problem behind this risk is in the **lack of control** after model release, such as the inability to monitor uses, the possibility to bypass the safeguards, and the fact that the public release of a model is irreversible (Eiras et al. 2024). The **lack of external accountability** is also a significant factor when examining the risks associated with model misuse. Developers contributing to open-source AI projects generally face very little external accountability for their contributions, even though they may feel responsible for the code they write and for preventing harms that directly arise from an implementation issue (Bartsch et al. 2024). At the same time, others may take refuge in the principle of neutrality inherent in open source software and feel that they are not responsible for any societal harm their work may have been used to cause (Widder et al. 2022). However, similar to many of the above-mentioned societal benefits, some of the risks relating to misuse are not yet substantiated by empirical evidence (Kapoor et al. 2024).

In addition, AI openness can make models more **vulnerable to security and privacy attacks** (Hintersdorf, Struppek, and Kersting 2025). Full access to the model makes "white box" attacks more feasible (Vassilev et al. 2024), such as adversarial example attacks (Ebrahimi et al. 2018; Sheikh and Zafar 2024), membership inference attacks (Nasr, Shokri, and Houmansadr 2019), and model inversion attacks (Chen et al. 2021). Admittedly, many of these attacks are also possible when the model is less "open" (e.g., through API access to the model), but local and free access to the model's weights make certain versions of these attacks possible that would otherwise be impossible or too expensive.

2.3 Discussion of AI in the Media

Since 2015, AI has been a topic regularly covered in French (Bellon and Velkovska 2023), American (Chuan, Tsai, and Cho 2019), and Chinese (Ittefaq et al. 2025) newspapers. The volume of articles on AI drastically increased in 2023 after the release of ChatGPT (Ittefaq et al. 2025; Tsimpoukis, Ratinaud, and Smyrnaios 2024). Previous work shows that Western newspaper articles published before 2020 tended to emphasize the benefits of AI more often than its risks (Chuan, Tsai, and Cho 2019; Cools, Van Gorp, and Opgenhaffen 2024). However, when an article is framed around the risks of AI—such as under a dystopian frame—those risks are discussed with more detailed examples and metaphors (Cools, Van Gorp, and Opgenhaffen 2024). Moreover, the proportion of critical articles has increased over time (Ittefaq et al. 2025; Nguyen and Hekman 2024). Articles specifically covering ethical issues are often neutral in tone and mainly address practical problems that may affect the public's daily life, such as privacy concerns (Ouchchy, Coin, and Dubljević 2020).

Previous research also shows that AI is portrayed very differently in the Western countries compared to China, where culture differs and the media is subject to more state censorship (Ji et al. 2024). In particular, Ittefaq et al. found Chinese newspapers to paint a highly positive picture of AI, emphasizing its benefits for the country's industrial development and economic growth (Ittefaq et al. 2025).

Our work follows the footsteps of these previous studies, particularly by analyzing similar elements (e.g., discussed risks and benefits), but focusing specifically on AI openness rather than AI more broadly.

3 Method

We focused on newspaper articles from the United States, France, and China because of the presence of leading companies that release open models in these countries (e.g., Meta's Llama 2 in the U.S., Mistral AI's Mistral 7B in France, and Deepseek's Deepseek R1 in China). Our study initially focused on the U.S. media. However, we observed that China and France were frequently mentioned in U.S. articles. We therefore included one journal from each country to diversify the perspectives. We limited our selection to newspapers written in French or English because the author conducting the analysis is proficient only in these languages. We acknowledge that this restricted our choice of Chinese newspapers. Within this language constraint, we chose newspapers based on their popularity. We therefore analyzed *The NYT*, *The WP*, *USA Today*, and *The WSJ* for the U.S., *Le Monde* for France, and *China Daily* for China. This leads to a choice of newspapers consistent with previous studies (Ittefaq et al. 2025; Chuan, Tsai, and Cho 2019; Cools, Van Gorp, and Opgenhaffen 2024). Given the rapid evolution in AI openness practices after 2020 (see Sec. 2.1), we chose to include articles published in or after 2020.

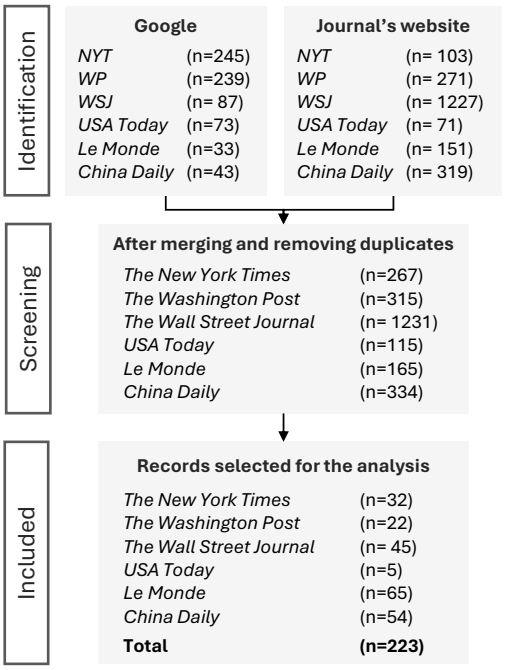

Figure 1: Our process of selecting newspaper articles.

3.1 Identification and Screening

The news articles were collected with Google search engine in incognito mode and with the search option of the news website. We queried articles from Jan. 1, 2020 to Mar. 31, 2025 with the following keywords: "open source ai"|"ai open source"|"open generative ai"|"ai openness"|"open ai"|"open source model"|"model openness"|"open model"|"open large language model"|"open source generative ai"|"open weights". To query *Le Monde*, we used French translations of those keywords.

After removing duplicates, we manually screened the articles for relevancy by reading the sentences where the term "open" is used (or "ouvert," "open," and "libre" for articles in French). We considered articles to be relevant if at least one sentence is entirely focused on AI openness (such as describing a benefit of AI openness or defining "open source AI"). A total of 223 articles were included from this process: 65 articles from *Le Monde*, 54 from *China Daily*, 5 from *USA Today*, 32 from the *NYT*, 45 from the *WSJ*, and 22 from the *WP* (See Fig. 1). The identification and screening process was performed from Feb. 15 to Apr. 19, 2025.

3.2 Analysis of the Articles

The first author, who is fluent in both French and English, started by reading all articles in full in their original written language to become familiar with the data. We then decided on a concrete analysis strategy for each article guided by the objective of this study, including two aspects: (1) examining the context in which AI openness is discussed—such as the focal topic of the article, the timeframe in which the article is published, and assessing the proportion of the article discussing AI openness—and (2) qualitatively analyzing how concretely the article discusses AI openness.

Context in which AI openness is discussed. We identified the focal topic of each article and grouped them based on the broader news category they fall under (e.g. model release, discussion on AI governance, etc.). For an article that shifts its focus from one to another (for example, starting with the release of a model and then moving on to discuss the broader AI market and competition), we prioritized the primary focus of the article as determined by the topic reflected in the title and lead paragraph.

We then identified the sections that address AI openness for each article. We classified articles based on the following categories: (1) The article is entirely dedicated to discussing AI openness; (2) A major portion of the article discusses AI openness, but not as the main subject; (3) A minor portion of the article discusses AI openness, but at least one full paragraph is focused on the topic; (4) Only a few isolated sentences of the article mention AI openness.

How AI openness is discussed. We used reflexive thematic analysis (Braun et al. 2022) to examine the sections identified as discussing AI openness. Once open coding was completed, we grouped the codes into themes that represent the different ways openness is addressed.

4 AI Openness in Newspapers

In this section, we first discuss how different articles situate the AI openness debate in specific contexts, such as the time of publication and the focal topic of the article (Sec. 4.1). Subsequently, we discuss how AI openness is discussed: what terminology is used, how the concept is defined, which AI models are cited as examples, how the open vs. closed debate is addressed and what arguments around AI openness are presented (Sec. 4.2).

4.1 Context in Which AI Openness Is Discussed

Openness in AI is a topic widely covered by newspapers in recent years: of the 223 articles included in our corpus 30 articles are entirely dedicated to the topic and 25 discuss it as a major part of the article. However, AI openness is not discussed uniformly across the newspapers, which not only address the topic with varying frequency and volume, but also through different topics.

Timeline. AI openness was barely discussed in newspapers before June 2023 (see Fig. 2), with only 7% (16/223) of the articles included in our analysis. The topic then started to garner increased interest throughout the second half of 2023 where several company decide to "open up" their models and received dedicated media coverage for it, such as MOSS (Shujuan 2023) and Qwen (Feifei 2023) in China, Llama 2 (Metz and Isaac 2023; Piquard 2023a; Isaac and Metz 2023; Vynck and Nix 2023) in the U.S., and Mistral in France (Metz 2023; Piquard 2023b). The topic continued to be regularly discussed in newspapers in 2024 with a decreased interest in late 2024 (respectively 75 and 59 articles were published in 2023 and 2024).

After the release of Deepseek R1 on 20 January 2025, however, the coverage of AI openness increased drastically,

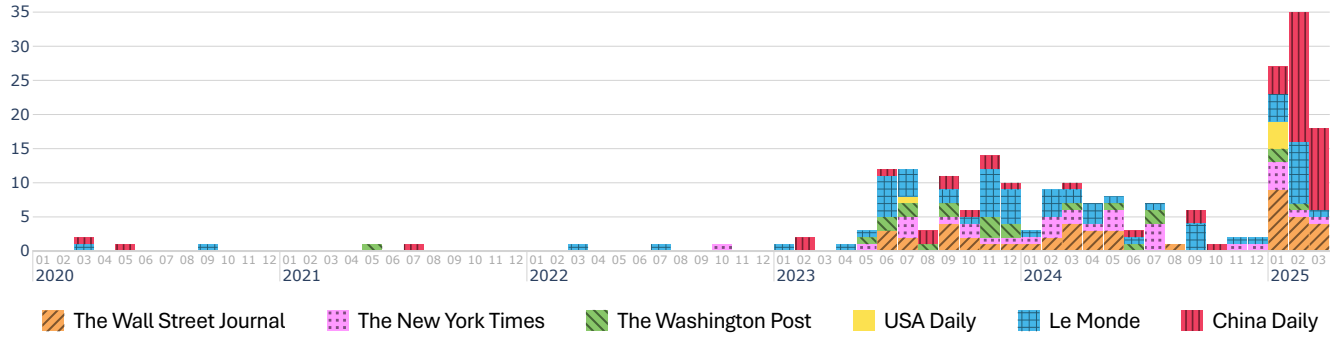

Figure 2: Distribution of the number of articles for each journal over time.

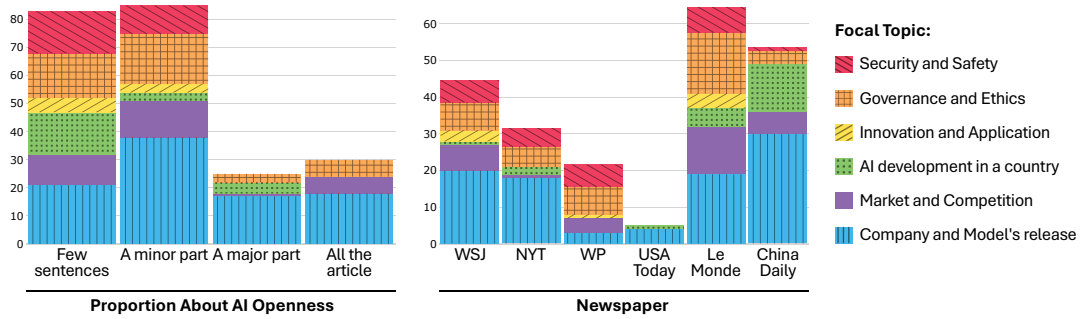

Figure 3: Distribution of articles and their focal topic based on the proportion of the article discussing openness (left) and the newspapers (right).

with 35% (79/223) of our corpus published since the release date, including 35 articles entirely dedicated to Deepseek. The spike in interest after Deepseek R1's release is more pronounced for *China Daily*, *USA Today*, and *the WSJ* (65% (35/54), 80% (4/5) and 40% (18/45) of the respective newspaper's articles were published after the release date) than for *The NYT*, *The WP*, and *Le Monde* (19% (6/32), 14% (3/22), and 20% (13/65), respectively).

Focal Topic. AI openness is most often discussed within a broader context and the open vs. closed debate is put into perspective alongside other current events. For example, openness practices are sometimes explained through an article focusing on harms caused by AI systems, while other times to present the new business model of a company. Identifying the focus of the article gives us crucial insight to understand the context in which AI openness is discussed and how it could shape the reader's opinion on the debate.

We grouped the articles into the following six focal topics: (1) *Security and Safety* (25/223), which raise alarms about the dangers of AI, often focusing on one particular security issue (e.g., bypassing a model's safeguards) or malicious application (e.g., deepfakes); (2) *Governance and Ethics* (43/223), presenting the debates and the diversity of perspectives on AI regulations, ethics, and policy activities (e.g., Paris AI Summit); (3) *Innovations and Applications* (8/223), which describes new AI techniques (e.g. distillation) or new AI applications (e.g. architecture); (4) *AI Development in a country* (22/223), which reports on government investments in AI (e.g., France building data centers) or AI development within a country or its regions; (5) *Market and Competition* (31/223), providing updates on the state of the "AI race", which companies are dominate the market, or current investment strategies in AI; (6) *Company and Model's Release* (94/223), which focuses on one particular model or company, often covering the release of a model and decisions made by this company (e.g., Meta releases Llama).

The *Company and Model Release* focal topic is not only the most common across the entire corpus, but also among the articles entirely dedicated to openness: 60% (18/30) of them fall within this theme, compared to 20% (6/30) within *Market and Competition*, and 20% (6/30) within *Governance and Ethics*. When the articles focus on *Innovation and Application* and *Security and Safety*, AI openness is most often discussed only as a small part of the article (See Figure 3).

The focal topic of articles discussing AI openness also varies significantly from one newspaper to another (see Fig. 3). For example, 56% (30/54) of the articles addressing AI openness in China Daily are focused on *Company and Model's Release* and 24% (13/54) on *AI Development in a country*, while other topics are much less covered. In contrast, *Governance and Ethics* and *Safety and Security* are more represented in other newspapers, such as the *WP*, where respectively 36% (8/22) and 27% (6/22) of the arti-

cles are focusing on these topics.

4.2 How AI Openness Is Discussed

Terminology. Although many terms are used in the articles to talk about openness in AI, "open source" is by far the most commonly used term, with $146/156$ of the English-language articles using it at least once to describe an AI system or a model. By comparison, the term "open weight(s)" is used in $4/156$ of the English-language articles, "open model" in $8/156$ and "openness" in $7/156$. In almost all articles written in French, the English term "open source" is used directly ($63/65$). These articles often accompany the English term with a range of different translations: "accès libre" ($18/65$), "libre accès" ($6/65$), and "source ouverte" ($1/65$). French-language articles also used the term "AI ouverte" ($10/65$) whereas its direct English translation "open AI (system)" is almost never used in English-language articles ($3/156$), likely to avoid confusion with the OpenAI company. Moreover, AI systems and models are sometimes described as "open source software" ($17/223$) and even "free software" (or "logiciel libre" in French) ($4/223$), even though the specific definitions of those terms and associated movements that do not fully apply to AI systems and their components (Open Source Initiative 2024; Paris, Moon, and Guo 2025; Sijbrandij 2023).

Definition. Across the articles, we identified 92 instances where openness in AI is defined: 19 in the *WSJ*, 29 in the *NYT*, 9 in the *WP*, 5 in *USA Today*, 23 in *Le Monde*, and 8 in *China Daily*. For each of these definitions, we analyzed its constituents, i.e., what makes AI "open" (see Tab. 1). Most definitions describe actions that can be performed when certain elements of the system is open.

A major difference between these definitions and those used in research (see Sec. 2.1) is that a large portion of definitions in the news article refer only to actions possible to be performed on the source code, without mentioning the model weights or the training data. While this may be due to the general public's and journalists' limited understanding of AI system structures, it also reflects a recurring confusion between open source software and openness in AI. One example is how Deepseek is described in the articles. Although the research article, model weights, and inference code of Deepseek R1 are available to the public, the Deepseek application on iOS and Android is proprietary. Nevertheless, one article reports:

> *"The DeepSeek app is completely open-source"* (Comstock 2025)

Models Described as "Open." Models the most often presented by newspapers as being "open" are (inclusive of any/all versions): Llama ($77/223$), Deepseek ($64/223$), Mistral ($37/223$), Qwen ($16/223$), Gemma ($10/223$), Stable Diffusion ($6/223$), Grok ($5/223$) and Falcon ($5/223$). Apart from Stable Diffusion, they are all large language models. They are considered to be "open" by the news articles because the model parameters are accessible. Notably, aside from Stable Diffusion and Falcon—both of which have partially accessible training data—all the other models reveal very little about their training datasets (Liesenfeld and Dingemanse 2024; Liesenfeld et al. 2025).

Category	Constituent of the definition	Count
What actions are enabled by openness?	Modify/Adapt	$40/92$
	Use	$32/92$
	Access	$32/92$
	Download/Copy	$14/92$
	Build upon	$10/92$
	Study/Inspect/View	$10/92$
	Share/Distribute	$10/92$
	Reproduce	$2/92$
Which part of the AI system do these actions apply to?	Source code	$35/92$
	Weights	$6/92$
	Technology	$5/92$
	Model	$5/92$
	Data	$4/92$
	System building blocks	$2/92$
	Software	$2/92$
	Algorithm	$1/92$
	Not specified	$40/92$
How these actions can be performed?	Freely/for free	$47/92$
	Without restriction/limits	$2/92$
	Without discrimination	$1/92$
Who can perform these actions?	Anyone	$33/92$
	Developers (only)	$6/92$
	Companies (only)	$5/92$
	Developers + Companies (only)	$3/92$
	Researchers + Companies (only)	$3/92$
	Developers + Government + Companies (only)	$1/92$
	Not specified	$41/92$
Other	"giving away"	$7/92$
	With a permissive/open source license	$4/92$

Table 1: Constituents making up the definitions of AI openness in newspapers and their number of occurrences.

Two news articles published in May and June 2023 (Metz and Isaac 2023; Larousserie 2023) refer to the first version of Llama released in February 2023 as "open source":

> *"Essentially, Meta was giving its A.I. technology away as open-source software—computer code that can be freely copied, modified and reused"* (Metz and Isaac 2023)

And even claim that Meta goes further than other open source AI project:

> *"But the company went further than many other open-source A.I. projects. It allowed people to download a version of LLaMA after it had been trained on enormous amounts of digital text culled from the internet."* (Metz and Isaac 2023)

In fact, the weights of this first version of Llama were not released under a license that allowed the same freedoms typically associated with open source, such as the ability to redistribute them. The access to the weights was granted by Meta only on a case-by-case basis to academics (Meta 2023). There are instances when the weights were leaked,

and Meta filed take-down requests to prevent the model weight from being redistributed (Github 2023; Cox 2023).

Openwashing and the Lack of Shared Definition. Although many articles talk about AI openness and define it as if it were an established concept, a few articles acknowledge the lack of a shared definition. For example, some articles mention that there is still no consensus on what "open source AI" means (Kessler 2024; Lin 2024a; Six 2025), or that the definition can vary depending of context (Lohr 2023). Some articles even acknowledge the efforts made to define and assess openness, such as those by the Linux Foundation with the Model Openness Framework (White et al. 2024; Kessler 2024), the OSI with its definition of Open Source AI (Open Source Initiative 2024; Kessler 2024; Six 2025), and the European Open Source AI Index (Liesenfeld et al. 2025; Six 2025).

Some also discuss the concept of openwashing, including an article from the NYT which defines it as follows:

> *"An accusation against some A.I. companies that they are using the 'open source' label too loosely."* (Kessler 2024)

Openwashing is often discussed in these news articles to call out companies frequently touting as contributing open source while not making their training data available, such as Meta, Mistral, and Deepseek.

Portrayal of the open vs. closed AI Debate. The discussion of whether AI systems should be open or closed is a major subject throughout our corpus. This is notably exemplified by having arguments for or against AI openness (e.g., stating that "the open approach has its critics too" (Vynck and Nix 2024)) as well as by explicitly describing the debate between advocates of openness and proponents of proprietary models. Newspapers present the open vs. closed debate as a central discussion that divides the AI community, describing it as "one of Silicon Valley's most divisive topics" (Lin 2024b), and stating that "Few debates have raged longer and more contentiously in the computing industry" (Isaac 2024). Some articles are entirely dedicated to explaining the debate (Isaac 2024; Mims 2025; Lin 2024b), and some stakeholders are publicly taking a stance on it through interviews (Piquard 2024b; Lohr 2023), opinion pieces (Lee 2023), and paid post (Meta 2024).

We found explicit mentions of such a debate in all journals except in China Daily, which presents a unique and consistent positive opinion on AI openness across its articles. However, this does not mean that the debate does not exist in China; it might happen in other media or platforms and can be indirectly observed in our corpus. In particular, an article from The NYT states:

> *"As in other countries, in China there is an intense debate over whether the latest technological advances should be made accessible to anyone or kept as closely held company secrets."* (Tobin and Metz 2024)

Arguments in Favor of or Against AI Openness. Most arguments in favor of or against AI openness presented in the articles are around the benefits and risks it presents for society. However, even if this dichotomy is widely covered in the

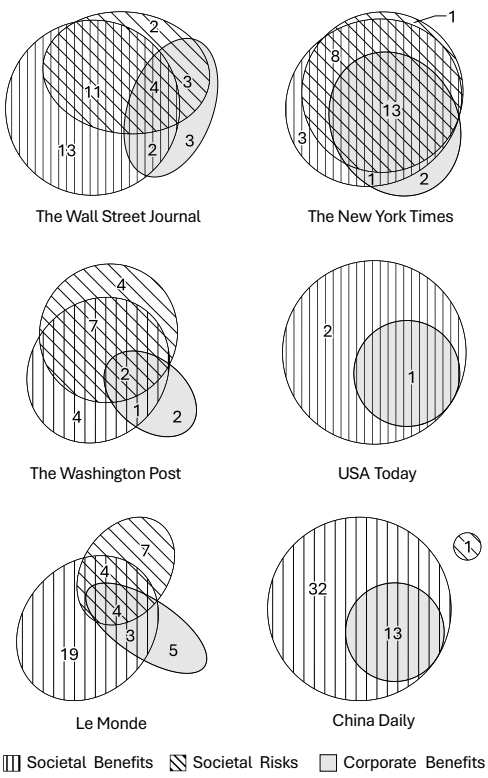

Figure 4: Distribution of articles discussing societal benefits, corporate benefits and societal risks of AI openness.

articles from our corpus, newspapers do not only weigh societal risks against societal benefits, but also contrast both factors with the benefits for companies in open-sourcing their models. Therefore, we categorize the argument on whether AI systems should be open or closed around the three factors: *societal benefits*, *societal risks* and *corporate benefits*.

Among our corpus, 66% $(147/223)$ of the articles mention at least one *societal benefit* of AI openness, 32% $(71/223)$ at least one *societal risk* and 26% $(59/223)$ at least one *corporate benefit*. The distribution of discussions on the benefits and risks of openness is highly uneven across journals (see Fig. 4). *China Daily* and *USA Today* focus significantly more on the benefits of openness than on the risks, whereas other journals take a more nuanced approach. However, we observe that even though *Le Monde* and the *WP* cover both societal risks and benefits, as well as corporate benefits, these three dimensions are more often separated into different articles than they are in *The New York Times*. Previous studies have shown that newspaper articles focusing solely on either benefits or risks tend to polarize public opinion more strongly (De Vreese and Boomgaarden 2006; Chong and Druckman 2007). This suggests that even if the overall frequency of coverage of risks and benefits is similar between those journals, the impact on the public may differ.

Societal Benefits. The newspaper articles mention a wide range of societal benefits, which we have organized into seven categories presented in Tab. 2. In addition to bene-

Category	Societal Benefits
Extensibility and Innovation (81/223)	Accelerate progress and innovation; Build new technology on top; More customizable; Enable experimentation; Enable application not possible without openness
Industry Development (55/223)	Reduce cost for downstream businesses; Lower the barriers of AI development; Boost competition; Serve global economic interests; Popularize AI
Autonomy (42/223)	A way around control of major firms; Provide more control; Ensure sovereignty; Avoid sharing data with third parties; Protect creative freedom; Cannot be taken over or derailed by institutions
Democratization and Inclusiveness (41/223)	Facilitates entrance of smaller players; Democratizes access; More inclusive; Avoid power concentration and monopoly; Shares opportunities and benefits; Democratizes AI development and research; Provides access to marginalized groups and regions; More equitable
Inspectability and Accountability (31/223)	Provide transparency; Permit to scrutinize and understand how the model is built; Improve trust; Allow public audit and tests; Enssure responsibility
Collaboration (29/223)	Fosters collaboration; Creates a community; Shares resources and knowledge
Security and safety (28/223)	More people can identify risks and bias; Makes AI safer; Makes the world safer

Table 2: Societal benefits associated with AI openness in newspapers (the number of articles mentioning at least one benefit of the category is indicated in parentheses).

fits that fall within these categories, 4 articles mention that openness is beneficial to society without specifying any particular benefits in the rest of the article.

Some of the societal benefits discussed in those articles are direct implications of a certain understanding of AI openness. For example, the ability to build new technologies on top of an open model is a direct consequence or direct quote from most AI openness definitions. On the other hand, other benefits discussed in some articles, such as making the world safer, are still debated among experts and do not represent a consensus in the AI community. These benefits may also depend on the types of model and its intended applications; for example, open sourcing a deepfake generative model will likely have more harmful consequences for society than a model designed to classify movies.

In terms of how these benefits are reported, 56% (178/318) of the quotations about benefits are framed as the opinion of a company or of "proponents of open-source AI models" (by quoting a person or attributing the statement to an organization) or as an ongoing debate that has not reached consensus in the AI community. Other benefits are framed as established truths or a consensus among experts. The benefits from the category *security and safety* are the one most often presented as subjective statements: 87% (39/45) of the quotes in this category are formulated as an someone's opinion. The benefits related to security and safety are themselves heavily questioned in newspaper articles (see Sec.4.2 *Societal Risks.*), which may partly explain why they are presented as more subjective than other benefits.

Category	Societal Risks
Misuse (44/223)	Spread disinformation and create propaganda campaigns; Permit to generate violent imagery, child pornography and non-consensual deepfakes; Facilitate discrimination, hate speech and harassment; Facilitate scams, cyberhackings and fishing campaigns; Help to create bomb, weapons and bioweapons; Privacy violation, especially using facial recognition
Lack of control (23/223)	Remove safeguards; Harder to track bad actors; Spread with no control; Irreversibility

Table 3: Societal risks associated with AI openness in newspapers (the number of articles mentioning at least one risk of the category is indicated in parentheses).

Moreover, some benefits reported in the newspapers depend on the meaning of "open source" AI. For example, some articles associate the ability to scrutinize the training data with AI openness, even if this capability is permitted by very few models called "open source" by the newspapers:

> *"OpenAI has been secretive about the data it uses to train its models, while some rivals are building open-source models that allow researchers to scrutinize the training data."* (Zakrzewski, Lima-Strong, and Oremus 2023)

Societal Risks. We identify two types of societal risks mentioned in newspapers: misuse of the models being open and the lack of control AI openness entails (see Tab. 3). In addition to risks that fall into these categories, 15 articles mention that openness is risky to society without specifying particular risks or harms it can cause in the rest of the article. Most of the articles consider the risks of misuse caused by the lack of control. For example, the generation of child pornography is mainly enabled by the possibility of removing safeguards from open models (Harwell 2023). Societal risks are presented more often as subjective claims: 64% (87/136) of the quotations about the risks are framed as a person's opinion or a non-consensual statement.

Corporate Benefits. The newspaper articles mention three type of corporate benefits: increasing market dominance, allowing to develop AI quicker and cheaper, and providing soft power (see Tab. 4). Corporate benefits are less often presented as subjective claims compared to societal benefits and risks: only 38% (41/108) of corporate benefit quotations are framed as subjective statements.

Unlike societal benefits and risks, corporate benefits can both be negatively and positively portrayed. Some articles present them as a positive aspect of openness, arguing that

Category	Corporate Benefits
Market dominance (47/223)	Boost adoption of the model; Upset the established dominance; Create an ecosystem; Not be locked in an other country's or competitor's ecosystem; Set standards; Have control over the future of AI
Quicker and Cheaper Development (27/223)	Benefit from independent developers; Speed up development; Catch up with other competitor; Recruitment tool
Soft power (8/223)	Gain respect and soft power

Table 4: Corporate benefits associated with AI openness in newspapers (the number of articles mentioning at least one benefit of the category is indicated in parentheses).

the growth of a company can benefit the public and the international community:

> *"DeepSeek's open-source community brings together developers from across the world who continuously improve their technologies or creations based on DeepSeek's core principles, facilitating the application and diffusion of technologies in a more flexible manner and on a wider scale."* (Dongmei 2025)

Other articles contrast them with societal risks or benefits, showing that the decision to open or close an AI model will ultimately be driven by personal gain rather than philanthropic motives:

> *"Open-weight models lack built-in security certifications, placing the burden of compliance on the deploying organization [...] But for DeepSeek, a goal may simply be to maximize its popularity and demonstrate to America's AI leaders that it has developed a powerful, cheaper technology."* (Lin 2025)

Finally, corporate benefits can be presented to highlight a sense of fatalism around AI openness: if a country does not support openness in its industry, its companies risk losing the "AI race," regardless of whether openness is socially good or bad:

> *"The United States also needs to resist stamping out open source AI. If the U.S. government does that, it risks ceding space to China and others."* (Cooper and James 2025)

These examples show that "open vs. closed AI" is not only an ethical dilemma, but also a debate about economics, international politics, and regulations.

5 Discussion

5.1 The Need for Terminological Clarity

Our analysis highlights that the news articles frequently present open source AI as a synonym of open source software. Some articles even used the term "free software," which have very different meaning and agenda from AI openness. The confusion between open source software and open source AI can be highly misleading for the public. As presented in Sec. 4.2, some articles incorrectly describe the proprietary software application built on top of a model as open. This can give the impression that an application is safer due to public transparency associated with openness when it is not. It can also let the reader believe that the benefits associated with open source software are transferable to AI systems, such as the ability to explain an output, which is less enabled by opening an AI system than by opening a software.

Moreover, the variance of definitions and meanings attributed to AI openness in the articles can add another layer to the disinformation. "Open source AI" is often used to describe models whose openness is still debated within the AI community (such as the first version of LLaMA that prevents redistribution or Llama 2 that restricts some uses (Maffulli 2023)). Some models are described as open in one article and closed in another, as is the case with GPT-2 in *Le Monde* (Piquard 2023c; Le Monde avec AFP 2024). The multitude of definitions and meanings can lead people to believe that some models provide some benefits that are not provided by these models, such as scrutinizing model's data.

All these examples highlight the urgency of defining consistent terminology surrounding AI openness and regulating the corresponding practices. While efforts have already been made—such as the multi-stakeholders process initiated by the Open Source Initiative to define open source AI (Open Source Initiative 2024)—these definitions still fall short of addressing the full complexity of open source issues in AI, which must encompass both the diversity of AI systems and the many nuances of AI openness.

5.2 Polarization of the Debate

News articles touch on most risks and benefits of AI openness discussed in academic research. However, they rarely mention the vulnerabilities to privacy and security attacks and the lack of external accountability. When describing these risks and benefits, reporters and interviewees quoted in the articles frequently use strong statements and hyperboles, as well as make broad generalizations. For example, whereas academic research suggests that openness could help identify bugs and safety issues in AI software more quickly (Seger et al. 2023), some newspaper articles claim that openness makes AI—or even the world—safer. In contrast, while academic research speaks of potential AI misuses and call for further empirical evidence (Kapoor et al. 2024), some narratives reported in newspapers assert that openness in AI could endanger society and compare opening AI system to opening the Manhattan Project (Mims 2025).

The polarization of the "open vs. closed AI" debate also reflects the ill-framing of the debate itself. This either-or framing not only obscures the many nuances of AI openness, but also conceals the diversity of AI as a technology and its applications. Very few articles in our corpus adapt their arguments based on the type of AI being discussed, its capabilities, or its domain of application. As in research, most discussions about openness are actually focused on generative AI and this narrow focus could lead to problematic generalizations that affect the regulations. Likewise, very few articles mention that there are many degrees of AI openness,

except in rare passages addressing openwashing. This binary framing is not solely conveyed by the journalists themselves; it often emerges from quotes from industry leaders, engineers, researchers, regulators, or government officials. Very few individuals interviewed in these articles express a nuanced opinion about AI openness. Moreover, when a company adopts a mixed approach by publishing both open and closed model, its commitment to the values of openness is called into question (Piquard 2024a).

Our results highlight the important role the AI community plays in shaping public understanding of AI governance practices such as AI openness. We need a more nuanced communication between the AI community and the public to counterbalance the current divisive narratives of AI openness. Arguments in favor and against openness in AI can coexist without necessarily splitting the community into two opposing camps.

5.3 Recommendation and Future Work

For the Journalists. We recommend the following practices to journalists and anyone wishing to communicate with the general public:

- Avoiding using the terms "open source software" and "free software" when referring to an AI system or model with publicly accessible components. The definitions of these terms, as established by organizations and in academic research, either do not apply to AI systems, or encompass a broader scope than just the accessibility of components. Acknowledging that openness in AI does not apply in the same way as it does to software can also help clarifying the difference.
- Preferring the term "open weights" to refer to the models that do not release their training data or code under non-restrictive license, such as LLaMA 2, LLaMA 3 and DeepSeek R1, even if these models are labeled as "open source models" by the companies. We recommend using the term "open source AI" either in accordance with the OSI definition (Open Source Initiative 2024), such as OLMo (Groeneveld et al. 2024), or for models that have released their weights, code (inference and training), data (training and fine-tuning), and documentation under an OSI-approved terms or OpenRAIL, such as BLOOM (Workshop et al. 2023). Although the terminology is not yet fully agreed upon within the AI community, distinguishing between "open weights" and "open source AI" is already a good starting point.
- Rely on public resources such as the European Open Source AI Index (Liesenfeld et al. 2025) and The Foundation Model Transparency Index (Bommasani et al. 2023) to verify the openness and transparency level of AI models.
- Specify the types of AI systems to which the risks and benefits of openness apply to avoid overgeneralization. Also acknowledge the different nuances of openness and the impact this has on both benefits and risks.
- Vary the individuals interviewed and quoted on the topic, particularly as there is an overrepresentation of Meta in American media and of Mistral in French media.

For the AI Community. For future research, we recommend the following directions:

- Explore how openness could apply to AI beyond generative AI and LLMs. In particular, envision definitions that could be consistent across a wide variety of AI systems and adaptable to future innovations. This is especially crucial for regulation and for ensuring consistent communication with the general public.
- Enrich and update public resources that assess the openness of each AI model. These resources are essential for effectively communicating the nuances of AI openness to the general public.
- Strengthen the understanding of the risks and benefits of AI openness based on empirical evidence. This could notably support public discourse by helping to confirm or refute claims made publicly by industry actors or government officials in the media.

5.4 Limitation

Although we diversified our news sources from three different countries, our analysis does not cover the entirety of the international press. In particular, press from countries less advanced in terms of AI innovation could present a very different perspective on AI openness. For example, we could hypothesize that these countries rely on open AI models produced by the U.S., France, and China, and that this may influence their approach to AI openness.

Even within the three selected countries, our results may not be representative of all media outlets. First, prior work have suggested that public perceptions of AI are influenced by the type of media consumed (Cui and Wu 2021). Since our study focuses on online newspapers, our findings may not be generalizable to all media formats. Second, we only selected a subset of newspapers published in each country, which means they do not represent the full political spectrum of each nation. Third, some articles relevant to our topic may not have been captured by our keywords. Finally, Chinese newspapers written in English may follow different editorial guidelines than those written in Chinese, given that they target a different audience. *China Daily* may also have an editorial line more closely shaped by the country's political agenda given that it is not an independent news outlet.

6 Conclusion

In this study, we analyzed how the media from the United States, China, and France portray openness practices in AI. In particular, we examined the contexts in which newspapers discuss AI openness, the meanings they assign to it and the arguments presented for and against the practice. This provides an overview of the messages the general public received through the media in different countries.

We find that certain media narratives on the topic is polarizing, and that the lack of clear and unified terminology about openness in AI leaves room for news articles to mislead readers, whether it be intentional or not. We therefore call on the AI community to work with journalists to reduce confusion and inaccuracies, and to recognize the importance of clarifying what openness should mean in AI.

Positionality Statement

Our perspectives are shaped by the countries where we grew up and spent significant amounts of time (including France, Korea, China, Canada, the U.S., and Switzerland), and the cultures specific to those places. The author who conducted the open coding stage of the reflexive thematic analysis regularly reads French and American media but is less familiar with Chinese media. Our academic training in computer science, robotics, and software engineering allows us to understand the technical details of AI systems. At the same time, our lack of experience in the field of journalism may limit our understanding of certain issues and experiences specific to the media industry.

Acknowledgments

We acknowledge the financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC). We thank Antoine Bou Khalil for his help and feedback, which contributed to the quality of this work.

References

Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Goodfellow, I.; Harp, A.; Irving, G.; Isard, M.; Jia, Y.; Jozefowicz, R.; Kaiser, L.; Kudlur, M.; Levenberg, J.; Mané, D.; Monga, R.; Moore, S.; Murray, D.; Olah, C.; Schuster, M.; Shlens, J.; Steiner, B.; Sutskever, I.; Talwar, K.; Tucker, P.; Vanhoucke, V.; Vasudevan, V.; Viégas, F.; Vinyals, O.; Warden, P.; Wattenberg, M.; Wicke, M.; Yu, Y.; and Zheng, X. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org.

Anderljung, M.; Hazell, J.; and von Knebel, M. 2025. Protecting society from AI misuse: when are restrictions on capabilities warranted? *AI & SOCIETY*, 40(5): 3841–3857.

Bartsch, S. C.; Lother, M.; Schmidt, J.-H.; Adam, M.; and Benlian, A. 2024. The Origin and Opportunities of Developers' Perceived Code Accountability in Open Source AI Software Development. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 7: 94–106.

Basdevant, A.; François, C.; Storchan, V.; Bankston, K.; Bdeir, A.; Behlendorf, B.; Debbah, M.; Kapoor, S.; LeCun, Y.; Surman, M.; King-Turvey, H.; Lambert, N.; Maffulli, S.; Marda, N.; Shivkumar, G.; and Tunney, J. 2024. Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence. arXiv:2405.15802.

Bellon, A.; and Velkovska, J. 2023. L'intelligence artificielle dans l'espace public : du domaine scientifique au problème public:Enquête sur un processus de publicisation controversé. *Réseaux*, 240(4): 31–70.

Bommasani, R.; Klyman, K.; Longpre, S.; Kapoor, S.; Maslej, N.; Xiong, B.; Zhang, D.; and Liang, P. 2023. The Foundation Model Transparency Index. arXiv:2310.12941.

Braun, V.; Clarke, V.; Hayfield, N.; Davey, L.; and Jenkinson, E. 2022. *Doing Reflexive Thematic Analysis*, 19–38. Cham: Springer International Publishing. ISBN 978-3-031-13942-0.

Brewer, P. R.; Bingaman, J.; Paintsil, A.; Wilson, D. C.; and Dawson, W. 2022. Media Use, Interpersonal Communication, and Attitudes Toward Artificial Intelligence. *Science Communication*, 44(5): 559–592.

Bucknall, B. S.; and Trager, R. F. 2023. Structured Access for Third-Party Research on Frontier AI Models: Investigating Researchers' Model Access Requirements. *AI Governance Initiative white paper*.

Chan, A.; Bucknall, B.; Bradley, H.; and Krueger, D. 2023. Hazards from Increasingly Accessible Fine-Tuning of Downloadable Foundation Models. In *Socially Responsible Language Modelling Research*.

Chen, S.; Kahla, M.; Jia, R.; and Qi, G.-J. 2021. Knowledge-Enriched Distributional Model Inversion Attacks. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 16158–16167.

Choi, S. 2024. Temporal Framing in Balanced News Coverage of Artificial Intelligence and Public Attitudes. *Mass Communication and Society*, 27(2): 384–405.

Choksi, M. Z.; Mandel, I.; and Benthall, S. 2025. The Brief and Wondrous Life of Open Models. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 3224–3240. New York, NY, USA: Association for Computing Machinery. ISBN 9798400714825.

Chong, D.; and Druckman, J. N. 2007. Framing Public Opinion in Competitive Democracies. *American Political Science Review*, 101(4): 637–655.

Chuan, C.-H.; Tsai, W.-H. S.; and Cho, S. Y. 2019. Framing Artificial Intelligence in American Newspapers. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 339–344. Honolulu HI USA: ACM. ISBN 978-1-4503-6324-2.

Claude AI Hub. 2024. Is Claude AI Open Source?

Clegg, N. 2024. Open Source AI Can Help America Lead in AI and Strengthen Global Security.

Comstock, L. 2025. What is DeepSeek? Why everyone's talking about this new AI. *USA TODAY*.

Cools, H.; Van Gorp, B.; and Opgenhaffen, M. 2024. Where exactly between utopia and dystopia? A framing analysis of AI and automation in US newspapers. *Journalism*, 25(1): 3–21.

Cooper, K. K.; and James. 2025. China's DeepSeek may erase America's lead in AI. And wreck your 401(k). | Opinion. *USA TODAY*.

Cox, J. 2023. Facebook's Powerful Large Language Model Leaks Online. *VICE*.

Cui, D.; and Wu, F. 2021. The influence of media use on public perceptions of artificial intelligence in China: Evidence from an online survey. *Information Development*, 37(1): 45–57.

De Vreese, C. H.; and Boomgaarden, H. G. 2006. Media Message Flows and Interpersonal Communication: The

Conditional Nature of Effects on Public Opinion. *Communication Research*, 33(1): 19–37.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.

Ding, J.; Akiki, C.; Jernite, Y.; Steele, A. L.; and Popo, T. 2023. Towards Openness Beyond Open Access: User Journeys through 3 Open AI Collaboratives. arXiv:2301.08488.

Dongmei, L. 2025. AI spread validates power of human progress. *China Daily*.

Ebrahimi, J.; Rao, A.; Lowd, D.; and Dou, D. 2018. HotFlip: White-Box Adversarial Examples for Text Classification. arXiv:1712.06751.

Eiras, F.; Petrov, A.; Vidgen, B.; Schroeder, C.; Pizzati, F.; Elkins, K.; Mukhopadhyay, S.; Bibi, A.; Purewal, A.; Botos, C.; Steibel, F.; Keshtkar, F.; Barez, F.; Smith, G.; Guadagni, G.; Chun, J.; Cabot, J.; Imperial, J.; Nolazco, J. A.; Landay, L.; Jackson, M.; Torr, P. H. S.; Darrell, T.; Lee, Y.; and Foerster, J. 2024. Risks and Opportunities of Open-Source Generative AI. arXiv:2405.08597.

European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA relevance). Legislative Body: CONSIL, EP.

Feifei, F. 2023. Alibaba doubles down on LLM race. *China Daily*.

Gao, K.; He, S.; He, Z.; Lin, J.; Pei, Q.; Shao, J.; and Zhang, W. 2023. Examining User-Friendly and Open-Sourced Large GPT Models: A Survey on Language, Multimodal, and Scientific GPT Models. arXiv:2308.14149.

Github. 2023. github/dmca - DMCA takedown notices 2023/03/2023-03-21-meta.

Groeneveld, D.; Beltagy, I.; Walsh, E.; Bhagia, A.; Kinney, R.; Tafjord, O.; Jha, A.; Ivison, H.; Magnusson, I.; Wang, Y.; Arora, S.; Atkinson, D.; Authur, R.; Chandu, K.; Cohan, A.; Dumas, J.; Elazar, Y.; Gu, Y.; Hessel, J.; Khot, T.; Merrill, W.; Morrison, J.; Muennighoff, N.; Naik, A.; Nam, C.; Peters, M.; Pyatkin, V.; Ravichander, A.; Schwenk, D.; Shah, S.; Smith, W.; Strubell, E.; Subramani, N.; Wortsman, M.; Dasigi, P.; Lambert, N.; Richardson, K.; Zettlemoyer, L.; Dodge, J.; Lo, K.; Soldaini, L.; Smith, N.; and Hajishirzi, H. 2024. OLMo: Accelerating the Science of Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 15789–15809. Bangkok, Thailand: Association for Computational Linguistics.

Harwell, D. 2023. AI-generated child sex images spawn new nightmare for the web. *Washington Post*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Hintersdorf, D.; Struppek, L.; and Kersting, K. 2025. Balancing Transparency and Risk: An Overview of the Security and Privacy Risks of Open-Source Machine Learning Models. In Steffen, B., ed., *Bridging the Gap Between AI and Reality*, 269–283. Cham: Springer Nature Switzerland. ISBN 978-3-031-73741-1.

Isaac, M. 2024. What to Know About the Open Versus Closed Software Debate. *The New York Times*.

Isaac, M.; and Metz, C. 2023. Meta Unveils a More Powerful A.I. and Isn't Fretting Over Who Uses It. *The New York Times*.

Ittefaq, M.; Zain, A.; Arif, R.; Ala-Uddin, M.; Ahmad, T.; and Iqbal, A. 2025. Global news media coverage of artificial intelligence (AI): A comparative analysis of frames, sentiments, and trends across 12 countries. *Telematics and Informatics*, 96: 102223.

Ji, X.; Kuai, J.; ; and Zamith, R. 2024. Scrutinizing Algorithms: Assessing Journalistic Role Performance in Chinese News Media's Coverage of Artificial Intelligence. *Journalism Practice*, 18(9): 2396–2413.

Kapoor, S.; Bommasani, R.; Klyman, K.; Longpre, S.; Ramaswami, A.; Cihon, P.; Hopkins, A.; Bankston, K.; Biderman, S.; Bogen, M.; Chowdhury, R.; Engler, A.; Henderson, P.; Jernite, Y.; Lazar, S.; Maffulli, S.; Nelson, A.; Pineau, J.; Skowron, A.; Song, D.; Storchan, V.; Zhang, D.; Ho, D. E.; Liang, P.; and Narayanan, A. 2024. On the Societal Impact of Open Foundation Models. arXiv:2403.07918.

Kembery, E.; Bucknall, B.; and Simpson, M. 2024. Position Paper: Model Access should be a Key Concern in AI Governance. arXiv:2412.00836.

Kessler, S. 2024. Openwashing. *The New York Times*.

Langenkamp, M. 2022. *How Open Source Machine Learning Software Shapes AI*. Thesis, Massachusetts Institute of Technology. Accepted: 2022-08-29T16:31:18Z.

Larousserie, D. 2023. Les concurrents de ChatGPT se multiplient. *Le Monde*.

Le Monde avec AFP. 2024. L'Unesco s'inquiète des préjugés sexistes de l'intelligence artificielle générative. *Le Monde*.

Lee, K.-F. 2023. Opinion | Artificial Intelligence Needs Open-Source Models to Reach Its Potential. *Wall Street Journal*.

Liesenfeld, A.; and Dingemanse, M. 2024. Rethinking open source generative AI: open washing and the EU AI Act. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1774–1787. Rio de Janeiro Brazil: ACM. ISBN 9798400704505.

Liesenfeld, A.; Dingemanse, M.; Blankvoort, D.; Kalra, N.; Kaya, A.; and Lopez, A. 2025. Open Source Generative AI Index: openness leaderboard.

Liesenfeld, A.; Lopez, A.; and Dingemanse, M. 2023. Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators. In *Proceedings of the 5th International Conference on Conversational User Interfaces*, CUI '23. New York, NY, USA: Association for Computing Machinery. ISBN 9798400700149.

Lin, B. 2024a. Open-Source Companies Are Sharing Their AI Free. Can They Crack OpenAI's Dominance? *Wall Street Journal*.

Lin, B. 2024b. Should AI Be Open-Source? Behind the Tweetstorm Over Its Dangers. *Wall Street Journal*.

Lin, B. 2025. How a Free Software Strategy Catapulted DeepSeek to AI Stardom. *Wall Street Journal*.

Lohr, S. 2023. An Industry Insider Drives an Open Alternative to Big Tech's A.I. *The New York Times*.

Maffulli, S. 2023. Meta's LLaMa 2 license is not Open Source.

Meta. 2023. Introducing LLaMA: A foundational, 65-billion-parameter language model.

Meta. 2024. As the AI race continues, open source models offer a path for all, not just the few. *The New York Times*.

Metz, C. 2023. Mistral, French A.I. Start-Up, Is Valued at $2 Billion in Funding Round. *The New York Times*.

Metz, C.; and Isaac, M. 2023. In Battle Over A.I., Meta Decides to Give Away Its Crown Jewels. *The New York Times*.

Mims, C. 2025. The Manhattan Project Was Secret. Should America's AI Work Be Too? *WSJ*.

Muñoz Ferrandis, C.; and Contractor, D. 2023. The BigScience OpenRAIL-M License.

Nasr, M.; Shokri, R.; and Houmansadr, A. 2019. Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, 739–753.

Nationen, V. 2024. *Govering AI for humanity: final report*. New York, NY: United Nations. ISBN 978-92-1-106787-3.

Nguyen, D.; and Hekman, E. 2024. The news framing of artificial intelligence: a critical exploration of how media discourses make sense of automation. *AI & SOCIETY*, 39(2): 437–451.

Nobel, P.; Rozenshtein, A. Z.; and Sharma, C. 2025. Unbundling AI Openness. In *WeRobot*.

Open Source Initiative. 2024. The Open Source AI Definition – 1.0.

Osborne, C. 2024. Why Companies "Democratise" Artificial Intelligence: The Case of Open Source Software Donations. arXiv:2409.17876.

Ouchchy, L.; Coin, A.; and Dubljević, V. 2020. AI in the headlines: the portrayal of the ethical issues of artificial intelligence in the media. *AI & SOCIETY*, 35(4): 927–936.

Palmer, A.; Smith, N. A.; and Spirling, A. 2024. Using proprietary language models in academic research requires explicit justification. *Nature Computational Science*, 4(1): 2–3.

Paris, T.; Moon, A.; and Guo, J. L. 2025. Opening the Scope of Openness in AI. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, 1293–1311. New York, NY, USA: Association for Computing Machinery. ISBN 9798400714825.

Perens, B.; et al. 1999. The open source definition. *Open sources: voices from the open source revolution*, 1: 171–188.

Piquard, A. 2023a. IA : Meta met la pression sur OpenAI et Google avec un modèle open source. *Le Monde*.

Piquard, A. 2023b. Mistral AI, incarnation des espoirs français dans l'intelligence artificielle. *Le Monde*.

Piquard, A. 2023c. OpenAI : aux sources de l'affaire Sam Altman. *Le Monde*.

Piquard, A. 2024a. Arthur Mensch, de Mistral AI : « En intelligence artificielle, nous pouvons rivaliser avec Google ou OpenAI ». *Le Monde*.

Piquard, A. 2024b. « Nous aimerions que tout le monde puisse créer sa propre IA ». *Le Monde*.

Seger, E.; Dreksler, N.; Moulange, R.; Dardaman, E.; Schuett, J.; Wei, K.; Winter, C.; Arnold, M.; hÉigeartaigh, S. Ó.; Korinek, A.; Anderljung, M.; Bucknall, B.; Chan, A.; Stafford, E.; Koessler, L.; Ovadya, A.; Garfinkel, B.; Bluemke, E.; Aird, M.; Levermore, P.; Hazell, J.; and Gupta, A. 2023. Open-Sourcing Highly Capable Foundation Models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. arXiv:2311.09227.

Sheikh, B. U. H.; and Zafar, A. 2024. White-box inference attack: compromising the security of deep learning-based COVID-19 diagnosis systems. *International Journal of Information Technology*, 16(3): 1475–1483.

Shevlane, T.; and Dafoe, A. 2020. The Offense-Defense Balance of Scientific Knowledge: Does Publishing AI Research Reduce Misuse? In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, 173–179. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-7110-0.

Shrestha, Y.; von Krogh, G.; and Feuerriegel, S. 2023. Building open-source AI. *Nature Computational Science*, 3(11): 908–911.

Shujuan, L. 2023. First domestic chatbot MOSS to be made open source. *China Daily*.

Sijbrandij, S. 2023. AI weights are not open "source".

Six, N. 2025. La part d'ombre des intelligences artificielles qui se disent « open source ». *Le Monde*.

Solaiman, I. 2023. The gradient of generative AI release: Methods and considerations. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*, 111–122.

Tobin, M.; and Metz, C. 2024. China Is Closing the A.I. Gap With the United States. *The New York Times*.

Tsimpoukis, P.; Ratinaud, P.; and Smyrnaios, N. 2024. Évolution des fréquences et des cooccurrences des entités nommées dans les discours de la presse sur l'intelligence artificielle (2012-2022). In *JADT 2024: Mots comptés, textes déchiffrés*, volume 2, 893–902. Bruxelles, Belgium: Presses universitaires de Louvain.

Vassilev, A.; Oprea, A.; Fordyce, A.; and Andersen, H. 2024. Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations. *NIST*. Last Modified: 2024-01-04T13:01-05:00.

Vynck, G. D.; and Nix, N. 2023. Facebook to make its AI free to use, expanding access to powerful tech. *Washington Post*.

Vynck, G. D.; and Nix, N. 2024. Meta releases open-source AI model it says rivals OpenAI, Google tech. *Washington Post*.

White, M.; Haddad, I.; Osborne, C.; Liu, X.-Y. Y.; Abdelmonsef, A.; Varghese, S.; and Hors, A. L. 2024. The Model Openness Framework: Promoting Completeness and Openness for Reproducibility, Transparency, and Usability in Artificial Intelligence. arXiv:2403.13784.

Whittlestone, J.; and Ovadya, A. 2020. The tension between openness and prudence in AI research. arXiv:1910.01170.

Widder, D. G.; Nafus, D.; Dabbish, L.; and Herbsleb, J. 2022. Limits and Possibilities for "Ethical AI" in Open Source: A Study of Deepfakes. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 2035–2046. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-9352-2.

Widder, D. G.; West, S.; and Whittaker, M. 2023. Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI. *Concentrated Power, and the Political Economy of Open AI*.

Workshop, B.; :; Scao, T. L.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; Gallé, M.; Tow, J.; Rush, A. M.; Biderman, S.; Webson, A.; Ammanamanchi, P. S.; Wang, T.; Sagot, B.; Muennighoff, N.; del Moral, A. V.; Ruwase, O.; Bawden, R.; Bekman, S.; McMillan-Major, A.; Beltagy, I.; Nguyen, H.; Saulnier, L.; Tan, S.; Suarez, P. O.; Sanh, V.; Laurençon, H.; Jernite, Y.; Launay, J.; Mitchell, M.; Raffel, C.; Gokaslan, A.; Simhi, A.; Soroa, A.; Aji, A. F.; Alfassy, A.; Rogers, A.; Nitzav, A. K.; Xu, C.; Mou, C.; Emezue, C.; Klamm, C.; Leong, C.; van Strien, D.; Adelani, D. I.; Radev, D.; Ponferrada, E. G.; Levkovizh, E.; Kim, E.; Natan, E. B.; Toni, F. D.; Dupont, G.; Kruszewski, G.; Pistilli, G.; Elsahar, H.; Benyamina, H.; Tran, H.; Yu, I.; Abdulmumin, I.; Johnson, I.; Gonzalez-Dios, I.; de la Rosa, J.; Chim, J.; Dodge, J.; Zhu, J.; Chang, J.; Frohberg, J.; Tobing, J.; Bhattacharjee, J.; Almubarak, K.; Chen, K.; Lo, K.; Werra, L. V.; Weber, L.; Phan, L.; allal, L. B.; Tanguy, L.; Dey, M.; Muñoz, M. R.; Masoud, M.; Grandury, M.; Šaško, M.; Huang, M.; Coavoux, M.; Singh, M.; Jiang, M. T.-J.; Vu, M. C.; Jauhar, M. A.; Ghaleb, M.; Subramani, N.; Kassner, N.; Khamis, N.; Nguyen, O.; Espejel, O.; de Gibert, O.; Villegas, P.; Henderson, P.; Colombo, P.; Amuok, P.; Lhoest, Q.; Harliman, R.; Bommasani, R.; López, R. L.; Ribeiro, R.; Osei, S.; Pyysalo, S.; Nagel, S.; Bose, S.; Muhammad, S. H.; Sharma, S.; Longpre, S.; Nikpoor, S.; Silberberg, S.; Pai, S.; Zink, S.; Torrent, T. T.; Schick, T.; Thrush, T.; Danchev, V.; Nikoulina, V.; Laippala, V.; Lepercq, V.; Prabhu, V.; Alyafeai, Z.; Talat, Z.; Raja, A.; Heinzerling, B.; Si, C.; Taşar, D. E.; Salesky, E.; Mielke, S. J.; Lee, W. Y.; Sharma, A.; Santilli, A.; Chaffin, A.; Stiegler, A.; Datta, D.; Szczechla, E.; Chhablani, G.; Wang, H.; Pandey, H.; Strobelt, H.; Fries, J. A.; Rozen, J.; Gao, L.; Sutawika, L.; Bari, M. S.; Al-shaibani, M. S.; Manica, M.; Nayak, N.; Teehan, R.; Albanie, S.; Shen, S.; Ben-David, S.; Bach, S. H.; Kim, T.; Bers, T.; Fevry, T.; Neeraj, T.; Thakker, U.; Raunak, V.; Tang, X.; Yong, Z.-X.; Sun, Z.; Brody, S.; Uri, Y.; Tojarieh, H.; Roberts, A.; Chung, H. W.; Tae, J.; Phang, J.; Press, O.; Li, C.; Narayanan, D.; Bourfoune, H.; Casper, J.; Rasley, J.; Ryabinin, M.; Mishra, M.; Zhang, M.; Shoeybi, M.; Peyrounette, M.; Patry, N.; Tazi, N.; Sanseviero, O.; von Platen, P.; Cornette, P.; Lavallée, P. F.; Lacroix, R.; Rajbhandari, S.; Gandhi, S.; Smith, S.; Requena, S.; Patil, S.; Dettmers, T.; Baruwa, A.; Singh, A.; Cheveleva, A.; Ligozat, A.-L.; Subramonian, A.; Névéol, A.; Lovering, C.; Garrette, D.; Tunuguntla, D.; Reiter, E.; Taktasheva, E.; Voloshina, E.; Bogdanov, E.; Winata, G. I.; Schoelkopf, H.; Kalo, J.-C.; Novikova, J.; Forde, J. Z.; Clive, J.; Kasai, J.; Kawamura, K.; Hazan, L.; Carpuat, M.; Clinciu, M.; Kim, N.; Cheng, N.; Serikov, O.; Antverg, O.; van der Wal, O.; Zhang, R.; Zhang, R.; Gehrmann, S.; Mirkin, S.; Pais, S.; Shavrina, T.; Scialom, T.; Yun, T.; Limisiewicz, T.; Rieser, V.; Protasov, V.; Mikhailov, V.; Pruksachatkun, Y.; Belinkov, Y.; Bamberger, Z.; Kasner, Z.; Rueda, A.; Pestana, A.; Feizpour, A.; Khan, A.; Faranak, A.; Santos, A.; Hevia, A.; Unldreaj, A.; Aghagol, A.; Abdollahi, A.; Tammour, A.; HajiHosseini, A.; Behroozi, B.; Ajibade, B.; Saxena, B.; Ferrandis, C. M.; McDuff, D.; Contractor, D.; Lansky, D.; David, D.; Kiela, D.; Nguyen, D. A.; Tan, E.; Baylor, E.; Ozoani, E.; Mirza, F.; Ononiwu, F.; Rezanejad, H.; Jones, H.; Bhattacharya, I.; Solaiman, I.; Sedenko, I.; Nejadgholi, I.; Passmore, J.; Seltzer, J.; Sanz, J. B.; Dutra, L.; Samagaio, M.; Elbadri, M.; Mieskes, M.; Gerchick, M.; Akinlolu, M.; McKenna, M.; Qiu, M.; Ghauri, M.; Burynok, M.; Abrar, N.; Rajani, N.; Elkott, N.; Fahmy, N.; Samuel, O.; An, R.; Kromann, R.; Hao, R.; Alizadeh, S.; Shubber, S.; Wang, S.; Roy, S.; Viguier, S.; Le, T.; Oyebade, T.; Le, T.; Yang, Y.; Nguyen, Z.; Kashyap, A. R.; Palasciano, A.; Callahan, A.; Shukla, A.; Miranda-Escalada, A.; Singh, A.; Beilharz, B.; Wang, B.; Brito, C.; Zhou, C.; Jain, C.; Xu, C.; Fourrier, C.; Periñán, D. L.; Molano, D.; Yu, D.; Manjavacas, E.; Barth, F.; Fuhrimann, F.; Altay, G.; Bayrak, G.; Burns, G.; Vrabec, H. U.; Bello, I.; Dash, I.; Kang, J.; Giorgi, J.; Golde, J.; Posada, J. D.; Sivaraman, K. R.; Bulchandani, L.; Liu, L.; Shinzato, L.; de Bykhovetz, M. H.; Takeuchi, M.; Pàmies, M.; Castillo, M. A.; Nezhurina, M.; Sänger, M.; Samwald, M.; Cullan, M.; Weinberg, M.; Wolf, M. D.; Mihaljcic, M.; Liu, M.; Freidank, M.; Kang, M.; Seelam, N.; Dahlberg, N.; Broad, N. M.; Muellner, N.; Fung, P.; Haller, P.; Chandrasekhar, R.; Eisenberg, R.; Martin, R.; Canalli, R.; Su, R.; Su, R.; Cahyawijaya, S.; Garda, S.; Deshmukh, S. S.; Mishra, S.; Kiblawi, S.; Ott, S.; Sang-aroonsiri, S.; Kumar, S.; Schweter, S.; Bharati, S.; Laud, T.; Gigant, T.; Kainuma, T.; Kusa, W.; Labrak, Y.; Bajaj, Y. S.; Venkatraman, Y.; Xu, Y.; Xu, Y.; Xu, Y.; Tan, Z.; Xie, Z.; Ye, Z.; Bras, M.; Belkada, Y.; and Wolf, T. 2023. BLOOM: A 176B-Parameter Open-Access Multilingual Language Model. arXiv:2211.05100.

Zakrzewski, C.; Lima-Strong, C.; and Oremus, W. 2023.

CEO behind ChatGPT warns Congress AI could cause 'harm to the world'. *Washington Post*.

Zhang, J.; Das, D.; Kamath, G.; and Tramer, F. 2025. Position: Membership Inference Attacks Cannot Prove That a Model was Trained on Your Data . In *2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 333–345. Los Alamitos, CA, USA: IEEE Computer Society.

Zuckerberg, M. 2024. Open Source AI Is the Path Forward.