

Should ChatGPT be biased? Challenges and risks of bias in large language models

by Emilio Ferrara

Abstract

As generative language models, exemplified by ChatGPT, continue to advance in their capabilities, the spotlight on biases inherent in these models intensifies. This paper delves into the distinctive challenges and risks associated with biases specifically in large-scale language models. We explore the origins of biases, stemming from factors such as training data, model specifications, algorithmic constraints, product design, and policy decisions. Our examination extends to the ethical implications arising from the unintended consequences of biased model outputs. In addition, we analyze the intricacies of mitigating biases, acknowledging the inevitable persistence of some biases, and consider the consequences of deploying these models across diverse applications, including virtual assistants, content generation, and chatbots. Finally, we provide an overview of current approaches for identifying, quantifying, and mitigating biases in language models, underscoring the need for a collaborative, multidisciplinary effort to craft AI systems that embody equity, transparency, and responsibility. This article aims to catalyze a thoughtful discourse within the AI community, prompting researchers and developers to consider the unique role of biases in the domain of generative language models and the ongoing quest for ethical AI.

Contents

- [1. Introduction](#)
- [2. Defining bias in generative language models](#)
- [3. Why are generative language models prone to bias?](#)
- [4. The inevitability of some forms of bias](#)
- [5. The broader risks of generative AI bias](#)
- [6. The role of human oversight and intervention](#)
- [7. Conclusions](#)

1. Introduction

In recent years, there has been a remarkable surge in the realm of artificial intelligence (AI), marked by the emergence of transformative technologies like ChatGPT and other generative language models (Brown, *et al.*, 2020; Devlin, *et al.*, 2019; OpenAI, 2023b; Ouyang, *et al.*, 2022; Radford, *et al.*, 2019, 2018; Schulman, *et al.*, 2017; Vaswani, *et al.*, 2017; Ziegler, *et al.*, 2019). These AI systems represent a class of sophisticated models designed to excel in generating human-like text and comprehending natural language. Using deep learning techniques and vast datasets, they have the ability to discern intricate patterns, make contextual

inferences, and generate coherent and contextually relevant responses to a diverse range of inputs (Bengio, *et al.*, 2000; Devlin, *et al.*, 2019; Mikolov, *et al.*, 2013; Sutskever, *et al.*, 2014; Vaswani, *et al.*, 2017). By mirroring human language capabilities, these models have unveiled a multitude of applications, ranging from chatbots and virtual assistants to translation services and content generation tools (Young, *et al.*, 2018).

Language models have ushered in a transformative era, underpinning the development of chatbots that emulate human interaction in conversations (Chen, *et al.*, 2017). These chatbots have become vital tools, simplifying customer service, technical support, and information queries with human-like interaction. Furthermore, the integration of language models into virtual assistants has endowed them with the ability to provide precise and contextually appropriate responses to user inquiries (S. Zhang, *et al.*, 2018). Virtual assistants, thus enhanced, become indispensable aides, capable of managing tasks ranging from appointment scheduling to Web searches and smart home device control (Wen, *et al.*, 2017).

Table 1: Factors contributing to bias in AI models.		
Contributing factor	Description	References
Training data	Biases in the source material or the selection process for training data can be absorbed by the model and reflected in its behavior.	Bender, <i>et al.</i> , 2021; Blodgett, <i>et al.</i> , 2020; Bolukbasi, <i>et al.</i> , 2016; Caliskan, <i>et al.</i> , 2017
Algorithms	Biases can be introduced or amplified through algorithms that place more importance on certain features or data points.	Blodgett, <i>et al.</i> , 2020; Hovy and Prabhumoye, 2021; Solaiman, <i>et al.</i> , 2019
Labeling and annotation	In (semi)supervised learning scenarios, biases may emerge from subjective judgments of human annotators providing labels or annotations for the training data.	Bender and Friedman, 2018; Buolamwini and Gebru, 2018; Munro, <i>et al.</i> , 2010
Product design decisions	Biases can arise from prioritizing certain use cases or designing user interfaces for specific demographics or industries, inadvertently	Benjamin, 2019; Kleinberg, <i>et al.</i> , 2016

	reinforcing existing biases and excluding different perspectives.	
Policy decisions	Developers might implement policies that prevent (or encourage) a given model behavior. For example, guardrails that modulate the behavior of ChatGPT and Bing-AI were designed to mitigate unintended toxic model behaviors or prevent malicious abuse.	Binns, 2018; Crawford, <i>et al.</i> , 2019; Doshi-Velez and Kim, 2017; Prates, <i>et al.</i> 2020

In the sphere of translation, harnessing the prowess of large language models facilitates markedly improved and fluent translations that span multiple languages, including those with limited resources (Costa-jussà, *et al.*, 2022; Karakanta, *et al.*, 2018; Pourdamghani and Knight, 2019; Ranathunga, *et al.*, 2023; Wang, *et al.*, 2019). Such capabilities not only foster enhanced cross-linguistic communication but can also enable timely solutions during emergencies and crises, especially in regions where low-resource languages or indigenous dialects are spoken (Christianson, *et al.*, 2018).

Moreover, the aptitude of language models to generate coherent and contextually pertinent text has rendered them invaluable in the realm of content creation. Acknowledged for their proficiency in producing various types of content, spanning articles, social media posts, and marketing materials, they have established a profound impact (Ferrara, 2023d; Yang, *et al.*, 2022).

These applications, among many others, underscore the transformative prowess of generative language models across an array of industries and sectors. However, as their adoption proliferates and their influence extends into ever-diverse domains (J.H. Choi, *et al.*, 2022; Gilson, *et al.*, 2023), it is imperative to confront the distinctive challenges posed by the potential biases that may be established within these models. These biases can have profound implications for users and society at large, highlighting the urgent need for comprehensive examination and mitigation of these issues (Ferrara, 2023c).

2. Defining bias in generative language models

2.1. Factors contributing to bias in large language models

Bias, in the context of large language models such as GPT-4 (Bubeck, *et al.*, 2023; OpenAI, 2023a, 2023b) and predecessors (Brown, *et al.*, 2020; Radford, *et al.*, 2019, 2018), or other state-of-the-art alternatives (Raffel, *et al.*, 2020; Touvron, *et al.*, 2023; including multimodal variants, Wu, *et al.*, 2023), can be defined as the presence of systematic misrepresentations, attribution errors, or factual distortions that result in favoring certain groups or ideas, perpetuating stereotypes, or making incorrect assumptions based on learned patterns. Biases in such models can arise due to several factors (*cf.*, [Table 1](#)).

One factor is the *training data*. If the data used to train a language model contain biases, either from the source material or through the selection process, these biases can be absorbed by the model and subsequently reflected in its behavior (Bender, *et al.*, 2021; Blodgett, *et al.*, 2020; Bolukbasi, *et al.*, 2016; Caliskan, *et al.*, 2017). Biases can also be introduced through the *algorithms* used to process and learn from the data. For example, if an algorithm places more importance on certain features or data points, it may unintentionally introduce or amplify biases present in the data (Blodgett, *et al.*, 2020; Hovy and Prabhunoye, 2021; Solaiman, *et al.*, 2019). In (semi)supervised learning scenarios, where *human annotators* provide labels or annotations for the training data, biases may emerge from the subjective judgments of the annotators themselves, influencing the model's understanding of the data (Bender and Friedman, 2018; Buolamwini and Gebru, 2018; Munro, *et al.*, 2010).

The choice of which *use cases* to prioritize or the design of user interfaces can also contribute to biases in large language models. For example, if a language model is primarily designed to generate content for a certain demographic or industry, it may inadvertently reinforce existing biases and exclude different perspectives (Benjamin, 2019; Kleinberg, *et al.*, 2016). Lastly, *policy decisions* can play a role in the manifestation of biases in language models. The developers of both commercial and openly available language models might implement policies that prevent (or encourage) a given model behavior. For example, both OpenAI and Microsoft have deliberate guardrails that modulate the behavior of ChatGPT and Bing-AI to mitigate unintended toxic model behaviors or prevent malicious abuse (Binns, 2018; Crawford, *et al.*, 2019; Doshi-Velez and Kim, 2017; Prate, *et al.*, 2020).

2.2. Types of biases in large language models

Large language models, which are commonly trained from vast amounts of text data present on the Internet, inevitably absorb the biases present in such data sources. These biases can take various forms (*cf.*, [Table 2](#)).

Table 2: Types of biases in large language models.		
Types of bias	Description	References
Demographic biases	These biases arise when the training data over-represents or under-represents certain demographic groups, leading the model to exhibit biased behavior towards specific genders, races, ethnicities, or other social groups.	Bender, <i>et al.</i> , 2021; Bolukbasi, <i>et al.</i> , 2016; Buolamwini and Gebru, 2018; Caliskan, <i>et al.</i> , 2017; Kirk, <i>et al.</i> , 2021; Munro, <i>et al.</i> , 2010
Cultural biases	Large language models may learn and perpetuate cultural stereotypes or biases, as they are often present in the data used for training. This	Blodgett, <i>et al.</i> , 2020; Bordia and Bowman, 2019;

	can result in the model producing outputs that reinforce or exacerbate existing cultural prejudices.	Ribeiro, <i>et al.</i> , 2020
Linguistic biases	Since the majority of the Internet's content is in English or a few other dominant languages, large language models tend to be more proficient in these languages. This can lead to biased performance and a lack of support for low-resource languages or minority dialects.	Bender, <i>et al.</i> , 2021; Conneau, <i>et al.</i> , 2017; Johnson, <i>et al.</i> , 2017; Pires, <i>et al.</i> , 2019; Ruder, <i>et al.</i> , 2019
Temporal biases	The training data for these models are typically restricted to limited time periods, or have temporal cutoffs, which may cause the model to be biased when reporting on current events, trends, and opinions. Similarly, the model's understanding of historical contexts or outdated information may be limited for lack of temporally representative data.	McCoy, <i>et al.</i> , 2019; Radford, <i>et al.</i> , 2018; N.A. Smith, 2020; Zellers, <i>et al.</i> , 2019
Confirmation biases	The training data may contain biases that result from individuals seeking out information that aligns with their pre-existing beliefs. Consequently, large language models may inadvertently reinforce these biases by providing outputs that confirm or support specific viewpoints.	Bolukbasi, <i>et al.</i> , 2016; Caliskan, <i>et al.</i> , 2017; Devlin, <i>et al.</i> , 2019; Mitchell, <i>et al.</i> , 2019
	Large language models can also learn and propagate the	

Ideological and political biases	political and ideological biases present in their training data. This can lead to the model generating outputs that favor certain political perspectives or ideologies, thereby amplifying existing biases.	Dixon, <i>et al.</i> , 2018; Garg, <i>et al.</i> , 2018; McCoy, <i>et al.</i> , 2019; McGee, 2023
----------------------------------	---	---

Demographic biases arise when the training data over-represents or under-represents certain demographic groups, leading the model to exhibit biased behavior towards specific genders, races, ethnicities, or other social groups (Bender, *et al.*, 2021; Bolukbasi, *et al.*, 2016; Buolamwini and Gebru, 2018; Caliskan, *et al.*, 2017; Kirk, *et al.*, 2021; Munro, *et al.*, 2010). Cultural biases occur when large language models learn and perpetuate cultural stereotypes or biases, as they are often present in the data used for training. This can result in the model producing outputs that reinforce or exacerbate existing cultural prejudices (Blodgett, *et al.*, 2020; Bordia and Bowman, 2019; Ribeiro, *et al.*, 2020). Linguistic biases emerge since the majority of the Internet's content is in English or a few other dominant languages, making large language models more proficient in these languages. This can lead to biased performance and a lack of support for low-resource languages or minority dialects (Bender, *et al.*, 2021; Conneau, *et al.*, 2017; Johnson, *et al.*, 2017; Pires, *et al.*, 2019; Ruder, *et al.*, 2019). Temporal biases appear as the training data for these models are typically restricted to limited time periods or have temporal cutoffs. This may cause the model to be biased when reporting on current events, trends, and opinions. Similarly, the model's understanding of historical contexts or outdated information may be limited due to a lack of temporally representative data (McCoy, *et al.*, 2019; Radford, *et al.*, 2018; N.A. Smith, 2020; Zellers, *et al.*, 2019). Confirmation biases in the training data may result from individuals seeking out information that aligns with their pre-existing beliefs. Consequently, large language models may inadvertently reinforce these biases by providing outputs that confirm or support specific viewpoints (Bolukbasi, *et al.*, 2016; Caliskan, *et al.*, 2017; Devlin, *et al.*, 2019; Mitchell, *et al.*, 2019). Lastly, ideological and political biases can be learned and propagated by large language models due to the presence of such biases in their training data. This can lead to the model generating outputs that favor certain political perspectives or ideologies, thereby amplifying existing biases (Dixon, *et al.*, 2018; Garg, *et al.*, 2018; McCoy, *et al.*, 2019; McGee, 2023).

This paper aims to explore the question of whether language models like GPT-4 (Bubeck, *et al.*, 2023; OpenAI, 2023a, 2023b), its prior versions (Brown, *et al.*, 2020; Radford, *et al.*, 2019, 2018), or other commercial or open source alternatives (Chowdhery, *et al.*, 2022; Raffel, *et al.*, 2020; Touvron, *et al.*, 2023) that power applications like ChatGPT (Ouyang, *et al.*, 2022) (or similar) should be biased or unbiased, taking into account the implications and risks of both perspectives. By examining the ethical, practical, and societal consequences of each viewpoint, we hope to contribute to the ongoing discussion surrounding responsible language model development and use. Through this exploration, our goal is to provide insights that can help guide the future evolution of GPT-style and other generative language models toward more ethical, fair, and beneficial outcomes while minimizing potential harm.

3. Why are generative language models prone to bias?

3.1. Biases from data

ChatGPT and other applications based on large language models are trained using a process that primarily relies on unsupervised learning, a machine learning technique that enables models to learn patterns and structures from vast amounts of unlabelled data (Carlini, *et al.*, 2021; Jiang, *et al.*, 2020). In most cases with these language models, the data consists of extensive text corpora available on the Internet, which includes Web sites, articles, books, and other forms of written content (Chowdhery, *et al.*, 2022; Devlin, *et al.*, 2019; Raffel, *et al.*, 2020; Touvron, *et al.*, 2023).

ChatGPT in particular is trained on a diverse range of Internet text datasets that encompass various domains, genres, and languages. While the specifics of the dataset used for GPT-4 are proprietary (OpenAI, 2023a, 2023b), the data sources utilized for training its predecessor, GPT-3, likely share similarities. For GPT-3 and predecessors, the primary dataset used was WebText (Radford, *et al.*, 2018), which is an ever-growing large-scale collection of Web pages (Brown, *et al.*, 2020; Radford, *et al.*, 2019). WebText was created by crawling the Internet and gathering text from Web pages. The sources of data include, but are not limited to:

- **Web sites:** Text is extracted from a wide array of Web sites, covering topics such as news, blogs, forums, and informational sites like Wikipedia. This enables the model to learn from diverse sources and gain knowledge on various subjects.
- **Books:** Text from books available online, including both fiction and non-fiction, contributes to the training data. This helps the model to learn different writing styles, narrative structures, and a wealth of knowledge from various fields.
- **Social media platforms:** Content from social media platforms, like Twitter, Facebook, and Reddit, is incorporated to expose the model to colloquial language, slang, and contemporary topics of discussion.
- **Conversational data:** To improve the model's conversational abilities, text from chat logs, comment sections, and other conversational sources are also included in the training dataset.

The developers of ChatGPT note that WebText data is preprocessed and filtered to remove low-quality content, explicit material, Web and social spam (Ferrara, 2022, 2019) and other undesirable text before being fed into the model (Brown, *et al.*, 2020; Radford, *et al.*, 2019). However, due to the vast scale of the data and the limitations of current filtering techniques, some undesirable or biased content may still seep into the training dataset, affecting the behavior of the resulting model. In addition to WebText, GPT-3 was further trained using a filtered version of the Common Crawl dataset (<https://commoncrawl.org>), a publicly available, massive Web-crawled dataset that contains raw Web page data, extracted metadata, and text content from billions of Web pages in multiple languages (Brown, *et al.*, 2020).

Another commonly-used dataset for language model training is *The Pile* (Biderman, *et al.*, 2022; Gao, *et al.*, 2020) an extensive and diverse collection of 22 smaller datasets, combining various sources of scientific articles, books, and Web content. It is designed for training large-scale language models, particularly in the domain of scientific research and understanding.

3.2. Biases from models

During the training process, generative language models are exposed to billions of sentences and phrases, allowing them to learn the intricate relationships between words, grammar, context, and meaning (Carlini, *et al.*, 2021; Jiang, *et al.*, 2020). As they process the text data, they gradually acquire natural language generation capabilities, enabling them to produce coherent and contextually relevant responses to various inputs. However, some capabilities of these models can lead to bias (*cf.*, [Table 3](#)):

Table 3: Sources of model bias in large language

models and their descriptions.		
Source	Description	References
Generalization	Models generalize knowledge from training data to new inputs, potentially leading to biased behavior if the data contains biases. This raises concerns about perpetuating biases, even if training data has been cleaned and filtered.	Caliskan, <i>et al.</i> , 2017; Gururangan, <i>et al.</i> , 2018
Propagation	Models may absorb and propagate biases in training data, adopting stereotypes and favoring certain groups or ideas, or making assumptions on non-representative learned patterns.	Bolukbasi, <i>et al.</i> , 2016; Dev and Phillips, 2019; Ferrara, 2023a
Emergence	Unanticipated capabilities and biases may emerge in large language models due to complex interactions between model parameters and biased training data. It has been proven difficult to predict or control these emergent biases.	Dettmers, <i>et al.</i> , 2022; Srivastava, <i>et al.</i> , 2022; Wei, Tay, <i>et al.</i> , 2022; Wei, Wang, <i>et al.</i> , 2022
Non-linearity	Biases in AI systems may have non-linear real-world impact, making it difficult to predict their consequences: small model biases may have massive negative effects, whereas large model biases might not cause significant consequences.	Chiappa, 2019; Ferrara, 2023a
	Reinforcement	

Alignment	Learning with Human Feedback (RLHF) finetunes large language models to reduce biases and align them with human values. The same principles might be abused to lead to unfair model behaviors.	Ouyang, <i>et al.</i> , 2022; Schulman, <i>et al.</i> , 2017; Ziegler, <i>et al.</i> , 2019
-----------	---	---

Generalization. One crucial aspect of these models is their ability to generalize, which allows them to apply the knowledge gained from their training data to new and previously unseen inputs, providing contextually relevant responses and predictions even in unfamiliar situations. However, this ability also raises concerns about potential biases, as models may inadvertently learn and perpetuate biases present in their training data, even if the data has been filtered and cleaned to the extent possible (Caliskan, *et al.*, 2017; Gururangan, *et al.*, 2018).

Propagation. As these models learn from the patterns and structures present in their training data, they may inadvertently absorb and propagate biases that they encounter, such as adopting stereotypes, favoring certain groups or ideas, or making assumptions based on learned patterns that do not accurately represent the full spectrum of human experience. This propagation of biases during training poses significant challenges to the development of fair and equitable AI systems, as biased models can lead to unfair treatment, reinforce stereotypes, and marginalize certain groups (Bolukbasi, *et al.*, 2016; Dev and Phillips, 2019; Ferrara, 2023a).

Emergence. In large language models, the phenomenon of emergence, which refers to the spontaneous appearance of unanticipated capabilities despite these functionalities not being explicitly encoded within the model's architecture or training data, can also result in unexpected biases due to the intricate interplay between model parameters and biased training data (Dettmers, *et al.*, 2022; Wei, Wang, *et al.*, 2022). The high-dimensional representations and non-linear interactions in these models make it difficult to predict or control these emergent biases, which may manifest in various ways, such as stereotyping, offensive language, or misinformation. To address this challenge, researchers are exploring bias mitigation strategies during training, fine-tuning with curated datasets, and *post hoc* emergent bias analyses (Srivastava, *et al.*, 2022; Wei, Tay, *et al.*, 2022).

Non-linearity. The non-linear relationships between biases in the system or data and their real-world impact imply that small biases may have massive negative effects, and large biases might not result in significant consequences. This disproportionality arises due to the complex interdependencies between the model parameters and the high-dimensional representations learned during training (Ferrara, 2023a). Randomized controlled trials could be used to draw causal relationships between the extent of each bias and their effects. In the absence of that, due to ethical reasons, multifaceted approaches involving in-depth analysis of model behavior, rigorous evaluation with diverse benchmarks, and the application of mitigation techniques that account for the nonlinear nature of emergent biases are needed (Chiappa, 2019).

Alignment. To address these issues, a strategy known as Reinforcement Learning with Human Feedback (RLHF) (Ziegler, *et al.*, 2019) was developed to fine-tune large language models like ChatGPT to reduce their biases and align them with human values. This approach involves collecting a dataset of human demonstrations, comparisons, and preferences to create a reward model that guides the fine-tuning process (Kumar, *et al.*, 2023). InstructGPT (ChatGPT's default model) (Ouyang, *et al.*, 2022) is trained using RLHF and then fine-tuned using Proximal Policy Optimization (PPO), a policy optimization algorithm

(Schulman, *et al.*, 2017). It is paramount to understand if the same principles could be exploited to deliberately misalign a model.

3.3. Can bias be mitigated with human-in-the-loop approaches?

Bias in generative language models can be mitigated to some extent with human-in-the-loop (HITL) approaches. These approaches involve incorporating human input, feedback, or oversight throughout the development and deployment of the language model, which can help address issues related to biases and other limitations. Here are some ways to integrate human-in-the-loop approaches to mitigate bias:

- **Training data curation:** Humans can be involved in curating and annotating high-quality and diverse training data. This may include identifying and correcting biases, ensuring a balance of perspectives, and reducing the influence of controversial or offensive content (Bender, *et al.*, 2021; Hovy and Prabhume, 2021; Hovy and Spruit, 2016).
- **Model fine-tuning:** Subject matter experts can guide the model fine-tuning process by providing feedback on the model's outputs, helping the model generalize better and avoid biased or incorrect responses (Gururangan, *et al.*, 2020).
- **Evaluation and feedback:** Human reviewers can evaluate the model's performance and provide feedback to developers, who can then iteratively improve the model. This feedback loop is essential for identifying and addressing bias-related issues. (Mitchell, *et al.*, 2019).
- **Real-time moderation:** Human moderators can monitor and review the model's outputs in real time, intervening when necessary to correct biased or inappropriate responses. This approach can be especially useful in high-stakes or sensitive applications (Park and Fung, 2017).
- **Customization and control:** Users can be provided with options to customize the model's behavior, adjusting the output according to their preferences or requirements. This approach can help users mitigate bias in the model's responses by tailoring it to specific contexts or domains (Bisk, *et al.*, 2020; Radford, *et al.*, 2018).

While human-in-the-loop approaches can help mitigate bias, it is essential to recognize that they may not be able to eliminate it entirely. Bias can stem from various sources, such as training data, fine-tuning process, or even the human reviewers themselves. However, combining machine learning techniques with human expertise can be a promising way to address some of the challenges posed by biases in generative language models.

4. The inevitability of some forms of bias

4.1. Are some biases inevitable?

Completely eliminating bias from large language models is a complex and challenging task due to the inherent nature of language and cultural norms. Since these models learn from vast amounts of text data available on the Internet, they are exposed to the biases present within human language and culture. Addressing bias in these models involves tackling several key challenges (*cf.*, [Table 4](#)).

Table 4: Challenges in addressing biases in large language models.

Challenge	Description	References
Inherent biases in language	Human language is a reflection of society, containing various biases, stereotypes, and assumptions. Separating useful patterns from these biases can be challenging as they are deeply ingrained in language structures and expressions.	Bourdieu, 1991; Fairclough, 2001; Foucault, 2002; Hill, 2008; Lakoff and Johnson, 1981; Whorf, 1964
Ambiguity of cultural norms	Cultural norms and values vary significantly across communities and regions. Determining which norms to encode in AI models is a complex task that requires a nuanced understanding of diverse cultural perspectives.	Geertz, 1973; Hofstede, 1980; Inglehart and Welzel, 2005; Triandis, 1995
Subjectivity of fairness	Fairness is a subjective concept with various interpretations. Eliminating bias from AI models requires defining “fair” in the context of applications, which is challenging due to the diverse range of stakeholders and perspectives.	Barocas, <i>et al.</i> , 2023; Friedler, <i>et al.</i> , 2016; Zafar, <i>et al.</i> , 2017
Continuously evolving language and culture	Language and culture constantly evolve, with new expressions, norms, and biases emerging over time. Keeping AI models up-to-date with these changes and ensuring they remain unbiased requires continuous monitoring and adaptation.	Castells, 2010; Jenkins and Deuze, 2008; Mufwene, 2001

First, human language is a reflection of society and as such, it contains various biases, stereotypes, and

assumptions. Separating useful patterns from these biases can be challenging, as they are often deeply ingrained in the way people express themselves and the structures of language itself (Bourdieu, 1991; Fairclough, 2001; Foucault, 2002; Hill, 2008; Lakoff and Johnson, 1981; Whorf, 1964). Second, cultural norms and values can vary significantly across different communities and regions. What is considered acceptable or appropriate in one context may be seen as biased or harmful in another. Determining which norms should be encoded in AI models and which should be filtered out is a complex task that requires careful consideration and a nuanced understanding of diverse cultural perspectives (Geertz, 1973; Hofstede, 1980; Inglehart and Welzel, 2005; Triandis, 1995).

Furthermore, fairness is a subjective concept that can be interpreted in various ways. Completely eliminating bias from AI models would require developers to define what “fair” means in the context of their applications, which can be a challenging task, given the diverse range of stakeholders and perspectives involved (Barocas, *et al.*, 2023; Friedler, *et al.*, 2016; Zafar, *et al.*, 2017). Lastly, language and culture are constantly evolving, with new expressions, norms, and biases emerging over time. Keeping AI models up-to-date with these changes and ensuring that they remain unbiased is an ongoing challenge that requires continuous monitoring and adaptation (Castells, 2010; Jenkins and Deuze, 2008; Mufwene, 2001).

Despite these challenges, it is essential for developers, researchers, and stakeholders to continue working towards reducing bias in large language models. By developing strategies for identifying and mitigating biases, collaborating with diverse communities, and engaging in ongoing evaluation and improvement, we can strive to create AI systems that are more equitable, fair, and beneficial for all users.

4.2. *Utility despite bias?*

Biased AI models can still be useful in certain contexts or applications, as long as users are aware of their limitations and take them into account when making decisions. In some cases, the biases present in these models may even be representative of the real-world context in which they are being used, providing valuable insights in surfacing societal inequalities that need to be tackled at their root.

The key to leveraging biased AI models responsibly is to ensure that users have a clear understanding of the potential biases and limitations associated with these models, so they can make informed decisions about whether and how to use them in different contexts. Some strategies for addressing this issue include:

- **Transparency:** Developers should be transparent about the methodologies, data sources, and potential biases of their AI models, providing users with the necessary information to understand the factors that may influence the model’s predictions and decisions. Best practices about the documentation of models and data have been advanced by the AI community (Gebru, *et al.*, 2021; Mitchell, *et al.*, 2019).
- **Education and awareness:** Providing resources, training, and support to help users better understand the potential biases in AI models and how to account for them when making decisions. This may involve creating guidelines, best practices, or other educational materials that explain the implications of bias in AI and how to navigate it responsibly.
- **Context-specific applications:** In some limited cases, biased AI models may be viable for specific applications or contexts where their biases align with the relevant factors or considerations. Experts should be employed to carefully evaluate the appropriateness of using biased models in these situations, taking into account the potential risks and benefits associated with their use, and actionable plans to recognize, quantify, and mitigate biases.
- **Continuous monitoring and evaluation:** Regularly assessing the performance of AI models in real-world contexts, monitoring their impact on users and affected communities, and making adjustments as needed to address any biases or unintended consequences that emerge over time (Bender, *et al.*, 2021).

By acknowledging that biased models can still be useful in limited contexts and taking steps to ensure that users are aware of their limitations and qualified to recognize and mitigate them, we can promote the responsible use of AI technologies and harness their potential benefits while minimizing the risks associated with bias.

5. The broader risks of generative AI bias

5.1. Pillars of responsible generative AI development

Ethical considerations of fairness and equality play a crucial role in the development and deployment of generative AI applications. As these models integrate increasingly into various aspects of our lives, their potential impact on individuals and society as a whole becomes a matter of significant concern. The responsibility lies with developers, researchers, and stakeholders to ensure that AI models treat all users and groups equitably, avoiding the perpetuation of existing biases or the creation of new ones (*cf.*, [Table 5](#)).

Table 5: Pillars of responsible generative AI development.		
Pillar	Description	References
Representation	Ensuring that the training data used to develop AI models is representative of the diverse range of perspectives, experiences, and backgrounds that exist within society. This helps to reduce the risk of biases being absorbed and propagated by the models, leading to more equitable outcomes.	Barocas, <i>et al.</i> , 2023; Buolamwini and Gebru, 2018; Mitchell, <i>et al.</i> , 2019; Sun, <i>et al.</i> , 2019; Torralba and Efros, 2011
Transparency	Developers should be transparent about the methodologies, data sources, and potential limitations of their AI models, enabling users to better understand the factors that may influence the model's predictions and decisions.	Ehsan, <i>et al.</i> , 2021; Larsson and Heintz, 2020

<p>Accountability</p>	<p>It is essential for developers and stakeholders to establish a clear framework for accountability, which may include monitoring the performance of AI models, addressing biases and errors, and responding to the concerns of users and affected communities.</p>	<p>Raji, <i>et al.</i>, 2020; H. Smith, 2021; Wachter, <i>et al.</i>, 2017</p>
<p>Inclusivity</p>	<p>AI applications should be designed to be inclusive and accessible to all users, taking into account factors such as language, culture, and accessibility needs, to ensure that the benefits of AI are shared equitably across society.</p>	<p>Morris, 2020; Schwartz, <i>et al.</i>, 2020</p>
<p>Protection of IP, human work, and human artistic expression</p>	<p>Generative AI models have the remarkable capability to create human-like text, artwork, music, and more. This creative aspect presents unique challenges, including issues related to intellectual property and the protection of human-generated copyright work to avoid AI plagiarism.</p>	<p>Ferrara, 2023c</p>
<p>Continuous improvement</p>	<p>Developers must commit to an ongoing process of evaluating, refining, and improving their AI models to address biases and ensure fairness over time. This may involve working with</p>	<p>Challen, <i>et al.</i>, 2019; Holstein, <i>et al.</i>, 2019; Mitchell, <i>et</i></p>

	researchers, policy-makers, and affected communities to gain information and feedback that can help guide the development of more equitable AI systems.	<i>al.</i> , 2019; O’Neil, 2016
--	---	---------------------------------

One key ethical consideration is *representation*. It is essential to ensure that the training data used to develop generative AI models are representative of the diverse range of perspectives, experiences, and backgrounds that exist within society (Barocas, *et al.*, 2023; Buolamwini and Gebru, 2018; Mitchell, *et al.*, 2019; Sun, *et al.*, 2019; Torralba and Efron, 2011). This helps to reduce the risk that biases are absorbed and propagated by models, leading to more equitable outcomes. *Transparency* is another important aspect. Developers should be transparent about the methodologies, data sources, and potential limitations of their generative AI models (Ehsan, *et al.*, 2021; Larsson and Heintz, 2020). This enables users to better understand the factors that may influence the model’s predictions and decisions. *Accountability* is also crucial for responsible generative AI development. Developers and stakeholders must establish a clear framework for accountability, which may include monitoring the performance of AI models, addressing biases and errors, and responding to users and communities affected (Raji, *et al.*, 2020; H. Smith, 2021; Wachter, *et al.*, 2017). An unique aspect of generative AI, compared to traditional machine learning, is its ability to possibly replace human artistic expression or to plagiarize the style and uniqueness of human work: as such, preservation of intellectual property, copyright protection and prevention of plagiarism are paramount (Ferrara, 2023c). *Inclusivity* is another key ethical consideration. Generative AI applications should be designed to be inclusive and accessible to all users, taking into account factors such as language, culture, and accessibility needs (Morris, 2020; Schwartz, *et al.*, 2020). This ensures that the benefits of AI are shared equitably across society.

Lastly, *continuous improvement* is vital to achieve fairness and equality in generative AI applications. Developers must commit to an ongoing process of evaluating, refining, and improving their AI models to address biases and ensure fairness over time (Challen, *et al.*, 2019; Holstein, *et al.*, 2019; Mitchell, *et al.*, 2019; O’Neil, 2016). This may involve collaborating with researchers, policy-makers, and affected communities to gain insight and feedback that can help guide the development of more equitable AI systems.

By prioritizing ethical considerations of fairness and equality, AI developers can create applications that not only harness the power of advanced technologies such as large language models, but also promote a more just and inclusive society, where the benefits and opportunities of AI are accessible to all.

5.2. The risks of exacerbating existing societal biases

Bias in widely-adopted AI models, including ChatGPT and other generative language models, can have far-reaching consequences that extend beyond the immediate context of their applications. When these models absorb and propagate biases, including those present in their training data, they may inadvertently reinforce stereotypes, marginalize certain groups, and lead to unfair treatment across various domains. Some examples of how biased AI models can adversely impact different areas include:

- **Hiring:** AI-driven hiring tools that use biased models may exhibit unfair treatment towards applicants from underrepresented groups or those with non-traditional backgrounds. This could lead to the perpetuation of existing inequalities in the job market, limiting opportunities for affected individuals and reducing diversity in the workforce (Bogen and Rieke, 2018; Raghavan, *et al.*, 2020). Large language models can be used to automate the screening of job applicants, such as by analyzing

resumes and cover letters. Since these models are trained on vast amounts of text data, they may have internalized biases present in the data, such as gender or racial biases. As a result, they could unintentionally favor certain applicants or disqualify others based on factors unrelated to their qualifications, reinforcing existing inequalities in the job market.

- **Lending:** Financial institutions increasingly rely on AI models for credit scoring and lending decisions. Biased models may unfairly penalize certain groups, such as minority communities or individuals with lower socio-economic status, by assigning them lower credit scores or denying them access to loans and financial services based on biased assumptions (Citron and Pasquale, 2014; N.T. Lee, *et al.*, 2019; Ustun, *et al.*, 2019). In lending, large language models can be used to assess creditworthiness or predict loan default risk, *e.g.*, based on automated analysis of application or support documents. If the data used to train these models contain historical biases or discriminatory lending practices, the models may learn to replicate these patterns. Consequently, they could deny loans to certain demographics or offer unfavorable terms based on factors like race, gender, or socioeconomic status, perpetuating financial inequality (Weidinger, *et al.*, 2021).
- **Content moderation:** AI-powered content moderation systems help manage and filter user-generated content on social media platforms and other online communities. If these systems are trained on biased data, they may disproportionately censor or suppress the voices of certain groups, while allowing harmful content or misinformation from other sources to proliferate (Augenstein, *et al.*, 2023; E.C. Choi and Ferrara, 2023; Gillespie, 2018; Roberts, 2019). Language models can be employed to automatically moderate and filter content on social media platforms or online forums. However, these models may struggle to understand the nuances of language, context, and cultural differences. They might over-moderate or under-moderate certain types of content, disproportionately affecting certain groups or topics. This could lead to censorship or the amplification of harmful content, perpetuating biases and misinformation (Davidson, *et al.*, 2019; Ezzeddine, *et al.*, 2023; Pasquetto, *et al.*, 2020; Sap, *et al.*, 2019).
- **Healthcare:** AI models are increasingly used to support medical decision-making and resource allocation. Biased models may result in unfair treatment for certain patient groups, leading to disparities in healthcare access and quality, and potentially exacerbating existing health inequalities (Challen, *et al.*, 2019; Gianfrancesco, *et al.*, 2018; Obermeyer, *et al.*, 2019; H. Smith, 2021). Large language models can be employed for tasks such as diagnosing diseases, recommending treatments, or analyzing patient data (Davenport and Kalakota, 2019; Ngiam and Khor, 2019). If the training data includes biased or unrepresentative information, the models may produce biased outcomes. Data used to train these models might be predominantly collected from specific populations, leading to less accurate predictions or recommendations for underrepresented groups. This can result in misdiagnoses, inadequate treatment plans, or unequal access to care. Models might unintentionally learn to associate certain diseases or conditions with specific demographic factors, perpetuating stereotypes and potentially influencing healthcare professionals' decision-making. Finally, biases in healthcare data could lead to models that prioritize certain types of treatments or interventions over others, disproportionately benefiting certain groups and disadvantaging others (Ferrara, 2023b; Paulus and Kent, 2020).
- **Education:** AI-driven educational tools and platforms can help personalize learning experiences and improve educational outcomes. However, biased models may perpetuate disparities by favoring certain learning styles or cultural backgrounds, disadvantaging students from underrepresented or marginalized communities (Elish and boyd, 2018; Selwyn, 2019). Large language models can be used in education for tasks such as personalized learning, grading, or content creation. If the models are trained on biased data, they may exacerbate existing biases in educational settings. Furthermore, if used for grading or assessing student work, language models might internalize biases in historical grading practices, leading to unfair evaluation of students based on factors like race, gender, or socioeconomic status. Finally, access in itself to ChatGPT or other AI tools for education can exacerbate preexisting inequalities.

5.3. Paths to AI transparency

Transparency and trust are essential components in the development and deployment of AI systems. As these models become more integrated into various aspects of our lives, it is increasingly important for users and regulators to understand how they make decisions and predictions, ensuring that they operate fairly, ethically, and responsibly.

Emphasizing transparency in AI systems can provide several benefits:

- **Informed decision-making:** When users and regulators have a clear understanding of how AI models make decisions and predictions, they can make more informed choices about whether to use or rely on these systems in different contexts. Transparency can empower users to evaluate the potential risks and benefits of AI systems and make decisions that align with their values and priorities (Goodman and Flaxman, 2017; Guidotti, *et al.*, 2018).
- **Public trust:** Fostering transparency can help build public trust in AI systems, as it demonstrates a commitment to ethical development and responsible deployment. When users and regulators can see that developers are taking steps to ensure the fairness and equity of their models, they may be more likely to trust and adopt these technologies (Ehsan, *et al.*, 2021; Larsson and Heintz, 2020).
- **Ethical compliance:** By promoting transparency, developers can demonstrate their compliance with ethical guidelines and regulations, showcasing their commitment to the responsible development of AI systems. This can help to establish a strong reputation and foster positive relationships with users, regulators, and other stakeholders (Jobin, *et al.*, 2019; Weidinger, *et al.*, 2021).
- **Collaborative improvement:** Transparency can facilitate collaboration between developers, researchers, policy-makers, and affected communities, enabling them to share insights and feedback that can help guide the development of more equitable and ethical AI systems (Challen, *et al.*, 2019; Holstein, *et al.*, 2019; Mitchell, *et al.*, 2019; O’Neil, 2016).

In summary, emphasizing transparency and trust in AI systems, including generative language models, is crucial for ensuring that users and regulators have a clear understanding of how these models make decisions and predictions. By promoting transparency, developers can demonstrate their commitment to ethical development and responsible deployment, fostering public trust and paving the way for more equitable and beneficial AI applications.

5.4. Regulatory efforts, industry standards, and ethical guidelines

As concerns about bias in AI systems continue to grow, several ongoing regulatory efforts and industry standards have emerged to address these challenges. AI ethics guidelines and fairness frameworks aim to provide guidance and best practices for developers, organizations, and policy-makers to reduce bias and ensure responsible development and deployment of AI systems (Floridi, 2019; Jobin, *et al.*, 2019).

Some notable efforts are summarized in [Table 6](#). These ongoing regulatory efforts and industry standards represent important steps in addressing bias and promoting the responsible development of AI systems. By adhering to these guidelines and frameworks, developers can contribute to the creation of AI technologies that are more equitable, fair, and beneficial for all users and communities.

Table 6: Ongoing regulatory efforts and industry standards to address bias in AI.

Effort	Description

<p>European Union's AI Ethics Guidelines</p>	<p>The EU's High-Level Expert Group on Artificial Intelligence has developed a set of AI Ethics Guidelines that outline key ethical principles and requirements for trustworthy AI, including fairness, transparency, accountability, and human agency (Independent High-Level Expert Group on Artificial Intelligence, 2019).</p>
<p>IEEE's Ethically Aligned Design</p>	<p>The Institute of Electrical and Electronics Engineers (IEEE) has published a comprehensive document, "Ethically Aligned Design," which provides recommendations and guidelines for the ethical development of AI and autonomous systems, with a focus on human rights, data agency, and technical robustness (Shahriari and Shahriari, 2017).</p>
<p>Partnership on AI</p>	<p>A coalition of tech companies, research institutes, and civil society organizations, the Partnership on AI aims to promote the responsible development and use of AI technologies. They work on various initiatives, including fairness, transparency, and accountability, to ensure that AI benefits all of humanity (Heer, 2018).</p>
<p>AI Fairness 360 (AIF360)</p>	<p>Developed by IBM Research, AIF360 is an open source toolkit that provides a comprehensive suite of metrics and algorithms to help developers detect and mitigate bias in their AI models. It assists developers in understanding and addressing fairness concerns in their AI applications (Bellamy, <i>et al.</i>, 2019).</p>
<p>Google's AI Principles</p>	<p>Google has outlined a set of AI Principles that guide the ethical development and use of AI technologies within the company. These principles emphasize fairness,</p>

	transparency, accountability, and the importance of avoiding harmful or unjust impacts (Pichai, 2018).
Algorithmic Impact Assessment (AIA)	Developed by the AI Now Institute, the AIA is a framework designed to help organizations evaluate and mitigate the potential risks and harms of AI systems. The AIA guides organizations through a structured process of identifying and addressing potential biases, discrimination, and other negative consequences of AI deployment (Reisman, <i>et al.</i> , 2018).
OECD Recommendation of the Council on Artificial Intelligence	The OECD Recommendation of the Council on Artificial Intelligence is a set of guidelines that provides a framework for the responsible development and deployment of AI, with five principles focused on inclusive growth, human-centered values, transparency and explainability, robustness, security and safety, and accountability (Cath, <i>et al.</i> , 2018; Yeung, 2020).

6. The role of human oversight and intervention

6.1. How to identify and mitigate bias?

Identifying and mitigating bias in AI models is essential for ensuring their responsible and equitable use. Various methods can be employed to address bias in AI systems:

- **Regular audits:** Conducting regular audits of AI models can help identify potential biases, errors, or unintended consequences in their outputs. These audits involve evaluating the model's performance against a set of predefined metrics and criteria, which may include fairness, accuracy, and representativeness. By monitoring AI models on an ongoing basis, developers can detect and address biases before they become problematic (Raji, *et al.*, 2020).
- **Retraining with curated data:** Retraining AI models with curated data can help reduce biases in their predictions and decisions. By carefully selecting and preparing training data that is more diverse, balanced, and representative of different perspectives, developers can ensure that AI models learn from a broader range of inputs and experiences, which may help mitigate the influence of biases present in the original training data (Gururangan, *et al.*, 2020; Ouyang, *et al.*, 2022).

- **Applying fairness metrics:** Fairness metrics can be used to evaluate the performance of AI models with respect to different user groups or populations (Ezzeldin, *et al.*, 2023; Yan, *et al.*, 2020). By analyzing AI model outputs based on these metrics, developers can identify potential disparities or biases in the model's treatment of different users and take steps to address them. Examples of fairness metrics include demographic parity, equalized odds, and equal opportunity (Mehrabi, *et al.*, 2021).
- **Algorithmic debiasing techniques:** Various algorithmic techniques have been developed to mitigate bias in AI models during training or post-processing. Some of these techniques include adversarial training, re-sampling, and re-weighting, which aim to minimize the influence of biased patterns and features on the model's predictions and decisions (Bender and Friedman, 2018; Bolukbasi, *et al.*, 2016; Bordia and Bowman, 2019; Dev and Phillips, 2019; N.T. Lee, *et al.*, 2019; Raghavan, *et al.*, 2020; Sun, *et al.*, 2019; B.H. Zhang, *et al.*, 2018).
- **Inclusion of diverse perspectives:** Ensuring that AI development teams are diverse and inclusive can help bring a wide range of perspectives and experiences to the table, which can contribute to the identification and mitigation of biases in AI models. By involving individuals from different backgrounds, cultures, and disciplines, developers can create more robust, fair, and representative AI systems (Biderman, *et al.*, 2022; Gao, *et al.*, 2020).
- **Human-in-the-loop approaches:** Incorporating human experts into the AI model development and decision-making processes can help provide valuable contextual understanding and ethical judgment that AI models may lack. Humans can serve as an additional layer of quality control, identifying biases, errors, or unintended consequences in AI model outputs and providing feedback to improve the model's performance and fairness.

Next, we dive deeper into the role that human experts can play in the responsible design and development of AI systems, including large language models, and their continuous oversight.

6.2. The importance of humans in the AI loop

Emphasizing the importance of involving human experts in AI system development, monitoring, and decision-making is crucial to ensuring the responsible and ethical deployment of AI technologies. Humans possess the ability to provide context and ethical judgment that AI models may lack, and their involvement can help address potential biases, errors, and unintended consequences that may arise from the use of these systems. Some key benefits of involving human experts in AI system development include:

- **Contextual understanding:** Human experts can provide valuable insights into the cultural, social, and historical contexts that shape language and communication, helping to guide AI models in generating more appropriate and sensitive responses (Dwivedi, *et al.*, 2021; Elish and boyd, 2018; Leslie, 2019).
- **Ethical judgment:** Human experts possess the moral and ethical reasoning skills needed to evaluate the potential impacts of AI systems on users and affected communities. By involving human experts in decision-making, we can ensure that AI models align with ethical principles and values, such as fairness, transparency, and accountability (Independent High-Level Expert Group on Artificial Intelligence, 2019; Jobin, *et al.*, 2019; Shahriari and Shahriari, 2017; Smuha, 2019; Weidinger, *et al.*, 2021).
- **Bias identification and mitigation:** Human experts can help identify and address biases in AI models, working alongside developers to implement strategies for mitigating or eliminating harmful biases and promoting more equitable and representative AI systems (Bender and Friedman, 2018; Raghavan, *et al.*, 2020; Sun, *et al.*, 2019).
- **Quality assurance and validation:** Human experts can serve as a vital layer of quality control,

evaluating AI model outputs for coherence, relevance, and potential biases, and providing feedback to improve the model's performance, accuracy, regulatory compliance, and trustworthiness (Felderer and Ramler, 2021).

- **Human override:** Incorporating human experts into AI system workflows can help strike a balance between automation and human judgment, allowing humans to intervene and override AI model decisions when necessary to ensure fairness, accountability, and ethical compliance (Etzioni and Etzioni, 2016).

By involving human experts in the development, monitoring, and decision-making processes of AI systems, we can leverage their contextual understanding and ethical judgment to complement the capabilities of AI models. This collaborative approach can help us create AI systems that are more responsible, equitable, and beneficial for all users, while also addressing the potential risks and challenges associated with bias in AI.

6.3. Possible strategies and best practices to address bias in generative AI

Addressing and mitigating potential biases in generative AI models requires a collaborative effort between AI developers, users, and affected communities. Fostering a more inclusive and fair AI ecosystem involves engaging various stakeholders in the development, evaluation, and deployment of AI technologies. This collaboration can oversee that AI models are designed to be more equitable, representative, and beneficial to all users.

Some key aspects of fostering collaboration in the AI ecosystem are shown in [Table 7](#). By fostering collaboration between AI developers, users, and affected communities, we can work towards creating a more inclusive and fair AI ecosystem that respects and values diverse perspectives and experiences. This collaborative approach can help ensure that AI technologies are developed and deployed in a way that is equitable, responsible, and beneficial for all users, while also addressing the potential risks and challenges associated with bias in AI.

Table 7: Strategies for addressing bias in generative AI systems.		
Strategy	Description	References
Engaging with affected communities	Involving affected communities in the development and evaluation of AI models can lead to the creation of generative AI systems that are more culturally sensitive, contextually relevant, and fair to all users.	Costanza-Chock, 2020; Eubanks, 2018; Gray and Suri, 2019; Mittelstadt, <i>et al.</i> , 2016; Taylor and Schroeder, 2015; Wallach, <i>et al.</i> , 2008; West, <i>et al.</i> , 2019
	Bringing together experts from different fields,	

Multidisciplinary collaboration	such as computer science, social sciences, humanities, and ethics, can help to develop more robust strategies for addressing and mitigating bias in generative AI systems.	Crawford, <i>et al.</i> , 2019; Holstein, <i>et al.</i> , 2019; Mittelstadt, <i>et al.</i> , 2016
User feedback and evaluation	Encouraging users to provide feedback on AI model outputs and performance can contribute to the ongoing improvement and refinement of generative AI models, ensuring that they remain fair, accurate, and relevant to users' needs.	Amershi, <i>et al.</i> , 2019; Cramer, <i>et al.</i> , 2008; M.K. Lee, <i>et al.</i> , 2015; Stoyanovich, <i>et al.</i> , 2020
Openness and transparency	Sharing information about the methodologies, data sources, and potential biases of generative AI models can enable stakeholders to make more informed decisions about whether and how to use these technologies in different contexts.	Cath, <i>et al.</i> , 2018; Jobin, <i>et al.</i> , 2019; Mittelstadt, <i>et al.</i> , 2016; Taddeo and Floridi, 2018
Establishing partnerships	Forming partnerships between AI developers, research institutions, non-profit organizations, and other stakeholders can facilitate the	Jobin, <i>et al.</i> , 2019

	sharing of knowledge, resources, and best practices, leading to the development of more equitable and responsible AI technologies.	
--	--	--



7. Conclusions

This paper highlights the challenges and risks associated with biases in generative language models like ChatGPT, emphasizing the need for a multi-disciplinary, collaborative effort to develop more equitable, transparent, and responsible AI systems that enhance a wide array of applications while minimizing unintended consequences.

Various methods for identifying and mitigating bias in AI models were presented, including regular audits, retraining with curated data, applying fairness metrics, and incorporating human experts in AI system development, monitoring, and decision-making. To achieve this goal, the development and deployment of AI technologies should prioritize ethical principles, such as fairness and equality, ensuring that all users and groups are treated equitably. Human oversight plays a vital role in providing context and ethical judgment that AI models may lack, helping to identify and address potential biases, errors, or unintended consequences. Collaboration between AI developers, users, and affected communities is essential for fostering a more inclusive and fair AI ecosystem, ensuring that diverse perspectives and experiences are considered and valued.

Continued research into methods for identifying, addressing, and mitigating biases in AI models will be critical to advancing the state of the art and promoting more equitable and inclusive AI systems. By bringing together experts from various disciplines, including computer science, social sciences, humanities, and ethics, we can foster a more comprehensive understanding of the potential biases and ethical challenges associated with AI applications.

Fostering an open and ongoing dialogue between stakeholders is crucial for sharing knowledge, best practices, and lessons learned from the development and use of AI applications. This dialogue can help to raise awareness of the potential risks and challenges associated with biases in AI models and promote the development of strategies and guidelines for mitigating their negative impacts.


Future research avenues. As the development of large language models continues, several essential aspects of research are necessary to advance their understanding and ensure their responsible deployment (*cf.*, [Table 8](#)). Among these are understanding their inner workings, addressing ethical concerns, ensuring controllability and safety, and developing more robust evaluation methods.

Table 8: Future avenues for large language models and generative AI research.	
Research	Description

areas	
Fairness, bias, and ethics	Addressing and minimizing biases in language models, as well as understanding their ethical implications, is a critical area of research; developing methods to detect, mitigate and prevent biases in AI models is essential.
Interpretability and explainability	Understanding the internal workings of large language models is a significant challenge; researchers are developing methods to make models more interpretable and explain their predictions.
Auditability and accountability	Large language models increasingly impact various sectors of society, influencing decision-making and shaping public discourse; ensuring that models are transparent, their actions can be traced, and those responsible for their development and deployment can be held accountable for the consequences of the AI's actions is vital for fostering trust and maintaining ethical standards in the AI community and beyond.
Controllable and safe AI	Ensuring that AI models can generate outputs that align with human intentions and values is an important research question; developing methods to control AI behavior, reduce harmful outputs and improve safety measures is vital.
Societal effects	The societal effects and implications of the deployment of AI systems encompass a wide range of concerns, including labor markets, privacy, bias, access to technology, public discourse, security, ethics, and regulation; Observing, characterizing, quantifying and understanding the broader effects that the deployment of AI systems has on society warrant careful consideration and continued research as these technologies continue to proliferate.

One critical area of research is fairness, bias, and ethics, which involves detecting, mitigating, and preventing biases in AI models to minimize their impact on various sectors of society. Another important research question is interpretability and explainability, as understanding the internal workings of large language models remains a significant challenge. Researchers are working to make models more interpretable and explain their predictions. Additionally, large language models can influence decision-making and shape public discourse, making auditability and accountability essential for fostering trust and maintaining ethical standards in the AI community and beyond.

Controllable and safe AI is also important to ensure that AI models can generate outputs that align with human intentions and values. Finally, observing, characterizing, quantifying, and understanding the broader effects that deploying AI systems has on society is critical for addressing societal effects, which encompass a wide range of concerns, including labor markets, privacy, bias, access to technology, public discourse, security, ethics, and regulation. The need for a multidisciplinary, collaborative effort to develop more equitable, transparent, and responsible AI systems is clear.

This paper emphasized the challenges and risks associated with biases in generative language models like ChatGPT and highlights the importance of continued research to develop responsible AI systems that enhance a wide array of applications while minimizing unintended consequences 

About the author

Emilio Ferrara is Professor at the University of Southern California, Research Team Leader at the USC Information Sciences Institute, and Principal Investigator at the USC/ISI Machine Intelligence and Data Science (MINDS).

E-mail: emiliofe [at] usc [dot] edu

Acknowledgements

The author is grateful to all current and past members of his lab at USC, and the numerous colleagues and students at USC Viterbi and Annenberg, who engaged in stimulating discussions and provided invaluable feedback about this study.

References

S. Amershi, D. Weld, M. Vorvoreanu, A. Fournery, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P.N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, and E.J. Horvitz, 2019. “Guidelines for human-AI interaction,” *CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, paper number 3, pp. 1–13.

doi: <https://doi.org/10.1145/3290605.3300233>, accessed 27 October 2023.

I. Augenstein, T. Baldwin, M. Cha, T. Chakraborty, G.L. Ciampaglia, D. Corney, R. DiResta, E. Ferrara, S. Hale, A. Halevy, E. Hovy, H. Ji, F. Menczer, R. Miguez, P. Nakov, D. Scheufele, S. Sharma, and G. Zagni, 2023. “Factuality challenges in the era of large language models,” *arXiv:2310.05189* (8 October).

doi: <https://doi.org/10.48550/arXiv.2310.05189>, accessed 27 October 2023.

S. Barocas, M. Hardt, and A. Narayanan, 2023. *Fairness and machine learning: Limitations and opportunities*. Cambridge, Mass.: MIT Press, and at <https://fairmlbook.org>, accessed 27 October 2023.

R.K. Bellamy, K., Dey, M. Hind, S.C. Hoffman, S. Houde, K. Kannan, P. Lohia, J. Martino, S. Mehta, and A. Mojsilović, 2019. “AI fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias,” *IBM Journal of Research and Development*, volume 63, numbers 4–5, pp. 4:1–4:15. doi: <https://doi.org/10.1145/10.1147/JRD.2019.2942287>, accessed 27 October 2023.

E.M. Bender and B. Friedman, 2018. “Data statements for natural language processing: Toward mitigating system bias and enabling better science,” *Transactions of the Association for Computational Linguistics*, volume 6, pp. 587–604, and at <https://aclanthology.org/Q18-1041/>, accessed 27 October 2023.

E.M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, 2021. “On the dangers of stochastic parrots: Can language models be too big?” *FACCT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623. doi: <https://doi.org/10.1145/3442188.3445922>, accessed 27 October 2023.

Y. Bengio, R. Ducharme, and P. Vincent, 2000. “A neural probabilistic language model,” *Advances in Neural Information Processing Systems*, volume 13, at https://papers.nips.cc/paper_files/paper/2000, accessed 27 October 2023.

R. Benjamin, 2019. *Race after technology: Abolitionist tools for the new jim code*. Medford, Mass.: Polity Press.

S. Biderman, K. Bicheno, and L. Gao, 2022. “Datasheet for the Pile,” *arXiv:2201.07311* (13 January). doi: <https://doi.org/10.48550/arXiv.2201.07311>, accessed 27 October 2023.

R. Binns, 2018. “Fairness in machine learning: Lessons from political philosophy,” *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81, pp. 149–159, and <https://proceedings.mlr.press/v81/binns18a.html>, accessed 27 October 2023.

Y. Bisk, A. Holtzman, J. Thomason, J. Andreas, Y. Bengio, J. Chai, M. Lapata, A. Lazaridou, J. May, A. Nisnevich, N. Pinto, and J. Turian, 2020. “Experience grounds language,” *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 8,718–8,735. doi: <https://doi.org/10.18653/v1/2020.emnlp-main.703>, accessed 27 October 2023.

S.L. Blodgett, S. Barocas, H. Daumé III, and H. Wallach, 2020. “Language (technology) is power: A critical survey of ‘bias’ in NLP,” *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5,454–5,476, and at <https://aclanthology.org/2020.acl-main.485/>, accessed 27 October 2023.

M. Bogen and A. Rieke, 2018. “Help wanted: An examination of hiring algorithms, equity, and bias,” *Upturn*, at <https://www.upturn.org/static/reports/2018/hiring-algorithms/files/Upturn%20--%20Help%20Wanted%20-%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf>, accessed 27 October 2023.

T. Bolukbasi, K.-W. Chang, J.Y. Zou, V. Saligrama, and A.T. Kalai, 2016. “Man is to computer programmer as woman is to homemaker? Debiasing word embeddings,” *NIPS'16: Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4,356–4,364.

S. Bordia and S.R. Bowman, 2019. “Identifying and reducing gender bias in word-level language models,” *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pp. 7–15, and at <https://aclanthology.org/N19-3002/>, accessed 27 October 2023.

P. Bourdieu, 1991. *Language and symbolic power*. Edited by J. Thompson. Translated by G. Raymond. Cambridge, Mass.: Harvard University Press.

T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, 2020. "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, volume 33, pp. 1,877–1,901, and at https://papers.nips.cc/paper_files/paper/2020, accessed 27 October 2023.

S. Bubeck, V. Chandrasekaran, R. Eldan, J. Gehrke, E. Horvitz, E. Kamar, P. Lee, Y.T. Lee, Y. Li, S. Lundberg, H. Nori, H. Palangi, M.T. Ribeiro, and Y. Zhang, 2023. "Sparks of artificial general intelligence: Early experiments with GPT-4," *arXiv:2303.12712* (22 March).
doi: <https://doi.org/10.48550/arXiv.2303.12712>, accessed 27 October 2023.

J. Buolamwini and T. Gebru, 2018. "Gender shades: Intersectional accuracy disparities in commercial gender classification," *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, pp. 77–91, and at <https://proceedings.mlr.press/v81/buolamwini18a.html>, accessed 27 October 2023.

A. Caliskan, J.J. Bryson, and A. Narayanan, 2017. "Semantics derived automatically from language corpora contain human-like biases," *Science*, volume 356, number 6334 (14 April), pp. 183–186.
doi: <https://doi.org/10.1126/science.aal4230>, accessed 27 October 2023.

N. Carlini, F. Tramer, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T.B. Brown, D. Song, U. Erlingsson, A. Oprea, and C. Raffel, 2021. "Extracting training data from large language models," *USENIX Security Symposium*, volume 6, at <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>, accessed 27 October 2023.

M. Castells, 2010. *The rise of the network society*. Malden, Mass.: Wiley-Blackwell.

C. Cath, S. Wachter, B. Mittelstadt, M. Taddeo, and L. Floridi, 2018. "Artificial intelligence and the 'good society': The US, EU, and UK approach," *Science and Engineering Ethics*, volume 24, pp. 505–528.
doi: <https://doi.org/10.1007/s11948-017-9901-7>, accessed 27 October 2023.

R. Challen, J. Denny, M. Pitt, L. Gompels, T. Edwards, and K. Tsaneva-Atanasova, 2019. "Artificial intelligence, bias and clinical safety," *BMJ Quality & Safety*, volume 28, number 3, pp. 231–237.
doi: <http://dx.doi.org/10.1136/bmjqs-2018-008370>, accessed 27 October 2023.

H. Chen, X. Liu, D. Yin, and J. Tang, 2017. "A survey on dialogue systems: Recent advances and new frontiers," *ACM SIGKDD Explorations Newsletter*, volume 19, number 2, pp. 25–35.
doi: <https://doi.org/10.1145/3166054.3166058>, accessed 27 October 2023.

S. Chiappa, 2019. "Path-specific counterfactual fairness," *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, number 01, pp. 7,801–7,808.
doi: <https://doi.org/10.1609/aaai.v33i01.33017801>, accessed 27 October 2023.

E.C. Choi and E. Ferrara, 2023. "Automated claim matching with large language models: Empowering fact-checkers in the fight against misinformation," *arXiv:2310.09223* (13 October).
doi: <https://doi.org/10.48550/arXiv.2310.09223>, accessed 27 October 2023.

J.H. Choi, K.E. Hickman, A.B. Monahan, and D. Schwarcz, 2022. "ChatGPT goes to law school," *Journal of Legal Education*, volume 71, number 3, pp. 387–400, and at <https://jle.aals.org/home/vol71/iss3/2/>, accessed 27 October 2023.

A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H.W. Chung, C. Sutton, S. Gehrmann, P. Schuh, K. Shi, S. Tsvyashchenko, J. Maynez, A. Rao, P. Barnes, Y. Tay, N.

Shazeer, V. Prabhakaran, E. Reif, N. Du, B. Hutchinson, R. Pope, J. Bradbury, J. Austin, M. Isard, G. Gur-Ari, P. Yin, T. Duke, A. Levskaya, S. Ghemawat, S. Dev, H. Michalewski, X. Garcia, V. Misra, K. Robinson, L. Fedus, D. Zhou, D. Ippolito, D. Luan, H. Lim, B. Zoph, A. Spiridonov, R. Sepassi, D. Dohan, S. Agrawal, M. Omernick, A.M. Dai, T.S. Pillai, M. Pellat, A. Lewkowycz, E. Moreira, R. Child, O. Polozov, K. Lee, Z. Zhou, X. Wang, B. Saeta, M. Diaz, O. Firat, M. Catasta, J. Wei, K. Meier-Hellstern, D. Eck, J. Dean, S. Petrov, and N. Fiedel, 2022. “PaLM: Scaling language modeling with pathways,” *arXiv:2204.02311* (5 April).

doi: <https://doi.org/10.48550/arXiv.2204.02311>, accessed 27 October 2023.

C. Christianson, J. Duncan, and B. Onyshkevych, 2018. “Overview of the DARPA LORELEI program,” *Machine Translation*, volume 32, pp. 3–9.

doi: <https://doi.org/10.1007/s10590-017-9212-4>, accessed 27 October 2023.

D.K. Citron and F. Pasquale, 2014. “The scored society: Due process for automated predictions,” *Washington Law Review*, volume 89, number 1, pp. 1–33, and at

<https://digitalcommons.law.uw.edu/wlr/vol89/iss1/2/>, accessed 27 October 2023.

A. Conneau, G. Lample, M. Ranzato, L. Denoyer, and H. Jégou, 2017. “Word translation without parallel data,” *arXiv:1710.04087* (11 October).

doi: <https://doi.org/10.48550/arXiv.1710.04087>, accessed 27 October 2023.

S. Costanza-Chock, 2020. *Design justice: Community-led practices to build the worlds we need*. Cambridge, Mass.: MIT Press.

doi: <https://doi.org/10.7551/mitpress/12255.001.0001>, accessed 27 October 2023.

M.R. Costa-jussà, J. Cross, O. Çelebi, M. Elbayad, K. Heafield, K. Heffernan, E. Kalbassi, J. Lam, D. Licht, J. Maillard, A. Sun, S. Wang, G. Wenzek, A. Youngblood, B. Akula, L. Barrault, G.M. Gonzalez, P. Hansanti, J. Hoffman, S. Jarrett, K.R. Sadagopan, D. Rowe, S. Spruit, C. Tran, P. Andrews, N.F. Ayan, S. Bhosale, S. Edunov, A. Fan, C. Gao, V. Goswami, F. Guzmán, P. Koehn, A. Mourachko, C. Ropers, S. Saleem, H. Schwenk, and J. Wang 2022. “No language left behind: Scaling human-centered machine translation,” *arXiv:2207.04672* (11 July).

doi: <https://doi.org/10.48550/arXiv.2207.04672>, accessed 27 October 2023.

H. Cramer, V. Evers, S. Ramlal, M. Van Someren, L. Rutledge, N. Stash, L. Aroyo, and B. Wielinga, 2008. “The effects of transparency on trust in and acceptance of a content-based art recommender,” *User Modeling and User-Adapted Interaction*, volume 18, pp. 455–496.

doi: <https://doi.org/10.1007/s11257-008-9051-3>, accessed 27 October 2023.

K. Crawford, R. Dobbe, T. Dryer, G. Fried, B. Green, E. Kaziunas, A. Kak, V. Mathur, E. McElroy, A.N. Sánchez, D. Raji, J.L. Rankin, R. Richardson, J. Schultz, S.M. West, and M. Whittaker, 2019. “AI Now 2019 report,” *AI Now Institute* (12 December), at <https://ainowinstitute.org/publication/ai-now-2019-report-2>, accessed 27 October 2023.

T. Davenport and R. Kalakota, 2019. “The potential for artificial intelligence in healthcare,” *Future Healthcare Journal*, volume 6, number 2, pp. 94–98.

doi: <https://doi.org/10.7861/futurehosp.6-2-94>, accessed 27 October 2023.

T. Davidson, D. Bhattacharya, and I. Weber, 2019. “Racial bias in hate speech and abusive language detection datasets,” *Proceedings of the Third Workshop on Abusive Language Online*, pp. 25–35.

doi: <https://doi.org/10.18653/v1/W19-3504>, accessed 27 October 2023.

T. Dettmers, M. Lewis, Y. Belkada, and L. Zettlemoyer, 2022. “LLM.int8(): 8-bit matrix multiplication for transformers at scale,” *arXiv:2208.07339* (15 August).

doi: <https://doi.org/10.48550/arXiv.2208.07339>, accessed 27 October 2023.

- S. Dev and J.M. Phillips, 2019. “Attenuating bias in word vectors,” *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, pp. 879–887, and at <https://proceedings.mlr.press/v89/dev19a.html>, accessed 27 October 2023.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, 2019. “BERT: Pre-training of deep bidirectional transformers for language understanding,” *Proceedings of NAACL-HLT 2019*, pp. 4,171–4,186, and at <https://aclanthology.org/N19-1423.pdf>, accessed 27 October 2023.
- L. Dixon, J. Li, J. Sorensen, N. Thain, and L. Vasserman, 2018. “Measuring and mitigating unintended bias in text classification,” *AIES '18: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 67–73.
doi: <https://doi.org/10.1145/3278721.3278729>, accessed 27 October 2023.
- F. Doshi-Velez and B. Kim, 2017. “Towards a rigorous science of interpretable machine learning,” *arXiv:1702.08608* (28 February).
doi: <https://doi.org/10.48550/arXiv.1702.08608>, accessed 27 October 2023.
- Y.K. Dwivedi, L. Hughes, E. Ismagilova, G. Aarts, C. Coombs, T. Crick, Y. Duan, R. Dwivedi, J. Edwards, A. Eirug, V. Galanos, P.V. Ilavarasan, M. Janssen, P. Jones, A.K. Kar, H. Kizgin, B. Kronemann, B. Lal, B. Lucini, R. Medaglia, and M.D. Williams, 2021. “Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy,” *International Journal of Information Management*, volume 57, 101994.
doi: <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>, accessed 27 October 2023.
- U. Ehsan, Q.V. Liao, M. Muller, M.O. Riedl, and J.D. Weisz, 2021. “Expanding explainability: Towards social transparency in AI systems,” *CHI '21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, article number 82, pp. 1–19.
doi: <https://doi.org/10.1145/3411764.3445188>, accessed 27 October 2023.
- M.C. Elish and d. boyd, 2018. “Situating methods in the magic of big data and AI,” *Communication Monographs*, volume 85, number 1, pp. 57–80.
doi: <https://doi.org/10.1080/03637751.2017.1375130>, accessed 27 October 2023.
- A. Etzioni and O. Etzioni, 2016. “Keeping AI legal,” *Vanderbilt Journal of Entertainment and Technology Law*, volume 19, number 1, pp. 133–146, at <https://scholarship.law.vanderbilt.edu/jetlaw/vol19/iss1/5/>, accessed 27 October 2023.
- V. Eubanks, 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. New York: St. Martin’s Press.
- F. Ezzeddine, O. Ayoub, S. Giordano, G. Nogara, I. Sbeity, E. Ferrara, and L. Luceri, 2023. “Exposing influence campaigns in the age of LLMs: A behavioral-based AI approach to detecting state-sponsored trolls,” *EPJ Data Science*, volume 12, number 1, article number 46.
doi: <https://doi.org/10.1140/epjds/s13688-023-00423-4>, accessed 27 October 2023.
- Y.H. Ezzeldin, S. Yan, C. He, E. Ferrara, and S. Avestimehr, 2023. “FairFed: Enabling group fairness in federated learning,” *AAAI'23/IAAI'23/EAAI'23: Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, article number 842, pp. 7,494–7,502.
doi: <https://doi.org/10.1609/aaai.v37i6.25911>, accessed 27 October 2023.
- N. Fairclough, 2001. *Language and power*. Second edition. New York: Longman.
- M. Felderer and R. Ramler, 2021. “Quality assurance for ai-based systems: Overview and challenges

(introduction to interactive session),” In: D. Winkler, S. Biffel, D. Mendez, M. Wimmer, and J. Bergsmann (editors). *Software quality: Future perspectives on software engineering quality*. Cham, Switzerland: Springer, pp. 33–42.

doi: https://doi.org/10.1007/978-3-030-65854-0_3, accessed 27 October 2023.

E. Ferrara, 2023a. “The butterfly effect in artificial intelligence systems: Implications for ai bias and fairness,” *arXiv:2307.05842* (11 July).

doi: <https://doi.org/10.48550/arXiv.2307.05842>, accessed 27 October 2023.

E. Ferrara, 2023b. Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies, *arXiv:2304.07683* (16 April).

doi: <https://doi.org/10.48550/arXiv.2304.07683>, accessed 27 October 2023.

E. Ferrara, 2023c. “GenAI against humanity: Nefarious applications of generative artificial intelligence and large language models,” *arXiv:2310.00737* (1 October).

doi: <https://doi.org/10.48550/arXiv.2310.00737>, accessed 27 October 2023.

E. Ferrara, 2023d. “Social bot detection in the age of ChatGPT: Challenges and opportunities,” *First Monday*, volume 28, number 6.

doi: <https://doi.org/10.5210/fm.v28i6.13185>, accessed 27 October 2023.

E. Ferrara, 2022. “Twitter spam and false accounts prevalence, detection, and characterization: A survey,” *First Monday*, volume 27, number 12.

doi: <https://doi.org/10.5210/fm.v27i12.12872>, accessed 27 October 2023.

E. Ferrara, 2019. “The history of digital spam,” *Communications of the ACM*, volume 62, number 8, pp. 82–91.

doi: <https://doi.org/10.1145/3299768>, accessed 27 October 2023.

L. Floridi, 2019. “Establishing the rules for building trustworthy AI,” *Nature Machine Intelligence*, volume 1, number 6 (7 May), pp. 261–262.

doi: <https://doi.org/10.1038/s42256-019-0055-y>, accessed 27 October 2023.

M. Foucault, 2002. *Archaeology of knowledge*. Second edition. London: Routledge.

S.A. Friedler, C. Scheidegger, and S. Venkatasubramanian, 2016. “On the (im)possibility of fairness,” *arXiv:1609.07236* (23 September).

doi: <https://doi.org/10.48550/arXiv.1609.07236>, accessed 27 October 2023.

L. Gao, S. Biderman, S. Black, L. Golding, T. Hoppe, C. Foster, J. Phang, H. He, A. Thite, N. Nabeshima, S. Presser, and C. Leahy, 2020. “The Pile: An 800GB dataset of diverse text for language modeling,” *arXiv:2101.00027* (31 December).

doi: <https://doi.org/10.48550/arXiv.2101.00027>, accessed 27 October 2023.

N. Garg, L. Schiebinger, D. Jurafsky, and J. Zou, 2018. “Word embeddings quantify 100 years of gender and ethnic stereotypes,” *Proceedings of the National Academy of Sciences*, volume 115, number 16 (3 April), pp. E3,635–E3,644.

doi: <https://doi.org/10.1073/pnas.1720347115>, accessed 27 October 2023.

T. Gebru, J. Morgenstern, B. Vecchione, J.W. Vaughan, H. Wallach, H.D. Iii, and K. Crawford, 2021. “Datasheets for datasets,” *Communications of the ACM*, volume 64, number 12, pp. 86–92.

doi: <https://doi.org/10.1145/3458723>, accessed 27 October 2023.

C. Geertz, 1973. *The interpretation of cultures: Selected essays*. New York: Basic Books.

M.A. Gianfrancesco, S. Tamang, J. Yazdany, and G. Schmajuk, 2018. "Potential biases in machine learning algorithms using electronic health record data," *JAMA Internal Medicine*, volume 178, number 11 (1 November), pp. 1,544–1,547.

doi: <https://doi.org/10.1001/jamainternmed.2018.3763>, accessed 27 October 2023.

T. Gillespie, 2018. *Custodians of the Internet: Platforms, content moderation, and the hidden decisions that shape social media*. New Haven, Conn.: Yale University Press.

A. Gilson, C.W. Safranek, T. Huang, V. Socrates, L. Chi, R.A. Taylor, and D. Chartash, 2023. "How does ChatGPT perform on the United States medical licensing examination? The implications of large language models for medical education and knowledge assessment," *JMIR Medical Education*, volume 9, number 1, e45312.

doi: <https://doi.org/10.2196/45312>, accessed 27 October 2023.

B. Goodman and S. Flaxman, 2017. "European Union regulations on algorithmic decision-making and a 'right to explanation'," *AI Magazine*, volume 38, number 3, pp. 50–57.

doi: <https://doi.org/10.1609/aimag.v38i3.2741>, accessed 27 October 2023.

M.L. Gray and S. Suri, 2019. *Ghost work: How to stop Silicon Valley from building a new global underclass*. Boston, Mass.: Houghton Mifflin Harcourt.

R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, 2018. "A survey of methods for explaining black box models" *ACM Computing Surveys*, volume 51, number 5, article number 93, pp. 1–42.

doi: <https://doi.org/10.1145/3236009>, accessed 27 October 2023.

S. Gururangan, A. Marasović, S. Swayamdipta, K. Lo, I. Beltagy, D. Downey, and N.A. Smith, 2020. "Don't stop pretraining: Adapt language models to domains and tasks," *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8,342–8,360.

doi: <https://doi.org/10.18653/v1/2020.acl-main.740>, accessed 27 October 2023.

S. Gururangan, S. Swayamdipta, O. Levy, R. Schwartz, S. Bowman, and N.A. Smith, 2018. "Annotation artifacts in natural language inference data," *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 2 (short papers), pp. 107–112, and at <https://aclanthology.org/N18-2017/>, accessed 27 October 2023.

J. Heer, 2018. "The partnership on AI," *AI Matters*, volume 4, number 3, pp. 25–26.

doi: <https://doi.org/10.1145/3284751.3284760>, accessed 27 October 2023.

J.H. Hill, 2008. *The everyday language of white racism*. Hoboken, N.J.: Wiley.

doi: <https://doi.org/10.1002/9781444304732>, accessed 27 October 2023.

G. Hofstede, 1980. *Culture's consequences: International differences in work-related values*. Beverly Hills, Calif.: Sage.

K. Holstein, J. Wortman Vaughan, H. Daumé III, M. Dudik, and H. Wallach, H. 2019. "Improving fairness in machine learning systems: What do industry practitioners need?" *CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, paper number 600, pp. 1–16.

doi: <https://doi.org/10.1145/3290605.3300830>, accessed 27 October 2023.

D. Hovy and S. Prabhunoye, 2021. "Five sources of bias in natural language processing," *Language and Linguistics Compass*, volume 15, number 8, e12432.

doi: <https://doi.org/10.1111/lnc3.12432>, accessed 27 October 2023.

D. Hovy and S.L. Spruit, 2016. "The social impact of natural language processing," *Proceedings of the 54th*

Annual Meeting of the Association for Computational Linguistics, volume 2: Short papers, pp. 591–598.
doi: <https://doi.org/10.18653/v1/P16-2096>, accessed 27 October 2023.

Independent High-Level Expert Group on Artificial Intelligence, 2019. “Ethics guidelines for trustworthy AI,” *European Commission* (8 April), at <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>, accessed 27 October 2023.

R. Inglehart and C. Welzel, 2005. *Modernization, cultural change, and democracy: The human development sequence*. Cambridge: Cambridge University Press.
doi: <https://doi.org/10.1017/CBO9780511790881>, accessed 27 October 2023.

H. Jenkins and M. Deuze, 2008. “Editorial: Convergence culture,” *Convergence*, volume 14, number 1, pp. 5–12.
doi: <https://doi.org/10.1177/1354856507084415>, accessed 27 October 2023.

Z. Jiang, F.F. Xu, J. Araki, and G. Neubig, 2020. “How can we know what language models know?” *Transactions of the Association for Computational Linguistics*, volume 8, pp. 423–438.
doi: https://doi.org/10.1162/tacl_a_00324, accessed 27 October 2023.

A. Jobin, M. Ienca, and E. Vayena, 2019. “The global landscape of AI ethics guidelines,” *Nature Machine Intelligence*, volume 1, number 9, pp. 389–399.
doi: <https://doi.org/10.1038/s42256-019-0088-2>, accessed 27 October 2023.

M. Johnson, M. Schuster, Q.V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado, M. Hughes, and J. Dean, 2017. “Google’s multilingual neural machine translation system: Enabling zero-shot translation,” *Transactions of the Association for Computational Linguistics*, volume 5, pp. 339–351, and at <https://aclanthology.org/Q17-1024/>, accessed 27 October 2023.

A. Karakanta, J. Dehdari, and J. van Genabith, 2018. “Neural machine translation for low-resource languages without parallel corpora,” *Machine Translation*, volume 32, pp. 167–189.
doi: <https://doi.org/10.1007/s10590-017-9203-5>, accessed 27 October 2023.

H.R. Kirk, Y. Jun, F. Volpin, H. Iqbal, E. Benussi, F.A. Dreyer, A. Shtedritski, and Y.M. Asano, 2021. “Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models,” *Advances in Neural Information Processing Systems*, volume 34, pp. 2,611–2,624, and at https://proceedings.neurips.cc/paper_files/paper/2021, accessed 27 October 2023.

J. Kleinberg, S. Mullainathan, and M. Raghavan, 2016. “Inherent trade-offs in the fair determination of risk scores,” *arXiv:1609.05807* (19 September).
doi: <https://doi.org/10.48550/arXiv.1609.05807>, accessed 27 October 2023.

V. Kumar, H. Koorehdavoudi, M. Moshtaghi, A. Misra, A. Chadha, and E. Ferrara, 2023. “Controlled text generation with hidden representation transformations,” *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 9,440–9,455.
doi: <https://doi.org/10.18653/v1/2023.findings-acl.602>, accessed 27 October 2023.

G. Lakoff and M. Johnson, 1981. *Metaphors we live by*. Chicago: University of Chicago Press.

S. Larsson and F. Heintz, 2020. “Transparency in artificial intelligence,” *Internet Policy Review*, volume 9, number 2.
doi: <https://doi.org/10.48550/10.14763/2020.2.1469>, accessed 27 October 2023.

M.K. Lee, D. Kusbit, E. Metsky, and L. Dabbish, 2015. “Working with machines: The impact of algorithmic and data-driven management on human workers,” *CHI ’15: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 1,603–1,612.

doi: <https://doi.org/10.1145/2702123.2702548>, accessed 27 October 2023.

N.T. Lee, P. Resnick, and G. Barton, 2019. “Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms,” *Brookings Institution* (22 May), at <https://www.brookings.edu/articles/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/>, accessed 27 October 2023.

D. Leslie, 2019. “Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector,” *Alan Turing Institute*.
doi: <https://doi.org/10.5281/zenodo.3240529>, accessed 27 October 2023.

R.T. McCoy, E. Pavlick, and T. Linzen, 2019. “Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference,” *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3,428–3,448, and at <https://aclanthology.org/P19-1334/>, accessed 27 October 2023.

R.W. McGee, 2023. “Is ChatGPT biased against conservatives? An empirical study” (17 February).
doi: <https://doi.org/10.2139/ssrn.4359405>, accessed 27 October 2023.

N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, 2021. “A survey on bias and fairness in machine learning,” *ACM Computing Surveys*, volume 54, number 6, article number 115, pp. 1–35.
doi: <https://doi.org/10.1145/3457607>, accessed 27 October 2023.

T. Mikolov, I. Sutskever, K. Chen, G.S. Corrado, and J. Dean, 2013. “Distributed representations of words and phrases and their compositionality,” *Advances in Neural Information Processing Systems*, volume 26, at https://papers.nips.cc/paper_files/paper/2013, accessed 27 October 2023.

M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I.D. Raji, and T. Gebru, 2019. “Model cards for model reporting,” *FAT* '19: Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 220–229.
doi: <https://doi.org/10.1145/3287560.3287596>, accessed 27 October 2023.

B.D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi, 2016. “The ethics of algorithms: Mapping the debate,” *Big Data & Society* (1 December).
doi: <https://doi.org/10.1177/2053951716679679>, accessed 27 October 2023.

M.R. Morris, 2020. “AI and accessibility,” *Communications of the ACM*, volume 63, number 6, pp. 35–37.
doi: <https://doi.org/10.1145/3356727>, accessed 27 October 2023.

S.S. Mufwene, 2001. *The ecology of language evolution*. New York: Cambridge University Press.
doi: <https://doi.org/10.1017/CBO9780511612862>, accessed 27 October 2023.

R. Munro, S. Bethard, V. Kuperman, V.T. Lai, R. Melnick, C. Potts, T. Schnoebelen, and H. Tily, 2010. “Crowdsourcing and language studies: The new generation of linguistic data,” *NAACL Workshop on Creating Speech and Language Data With Amazon’s Mechanical Turk*, pp. 122–130, and at <https://aclanthology.org/W10-0719/>, accessed 27 October 2023.

K.Y. Ngiam and I.W. Khor, 2019. “Big data and machine learning algorithms for health-care delivery,” *Lancet Oncology*, volume 20, number 5, pp. e262–e273.
doi: [https://doi.org/10.1016/S1470-2045\(19\)30149-4](https://doi.org/10.1016/S1470-2045(19)30149-4), accessed 27 October 2023.

Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, 2019. “Dissecting racial bias in an algorithm used to manage the health of populations,” *Science*, volume 366, number 6464 (25 October), pp. 447–453.
doi: <https://doi.org/10.1126/science.aax2342>, accessed 27 October 2023.

C. O’Neil, 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Crown.

OpenAI, 2023a. “GPT-4 System Card” (23 March), at <https://cdn.openai.com/papers/gpt-4-system-card.pdf>, accessed 27 October 2023.

OpenAI, 2023b. “GPT-4 technical report” (27 March), at <https://cdn.openai.com/papers/gpt-4.pdf>, accessed 27 October 2023.

L. Ouyang, J. Wu, X. Jiang, D. Almeida, C.L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, J. Schulman, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. Christiano, J. Leike, and R. Lowe, 2022. “Training language models to follow instructions with human feedback,” *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, at https://proceedings.neurips.cc/paper_files/paper/2022, accessed 27 October 2023.

J.H. Park and P. Fung, 2017. “One-step and two-step classification for abusive language detection on Twitter,” *Proceedings of the First Workshop on Abusive Language Online*, pp. 41–45. doi: <https://doi.org/10.18653/v1/W17-3006>, accessed 27 October 2023.

I.V. Pasquetto, B. Swire-Thompson, M.A. Amazeen, F. Benevenuto, N.M. Brashier, R.M. Bond, L.C. Bozarth, C. Budak, U.K. Ecker, L.K., Fazio, E. Ferrara, A.J. Flanagin, A. Flammini, D. Freelon, N. Grinberg, R. Hertwig, K.H. Jamieson, K. Joseph, J.J. Jones, R.K. Garrett, D. Kreiss, S. McGregor, J. McNealy, D. Margolin, A. Marwick, F. Menczer, M.J. Metzger, S. Nah, S. Lewandowsky, P. Lorenz-Spreen, P. Ortellado, G. Pennycook, E. Porter, D.G. Rand, R.E. Robertson, F. Tripodi, S. Vosoughi, C. Vargo, O. Varol, B.E. Weeks, J. Wihbey, T.J. Wood, and K.-C. Yang, 2020. “Tackling misinformation: What researchers could do with social media data,” *Harvard Kennedy School Misinformation Review* (9 December), at <https://misinforeview.hks.harvard.edu/article/tackling-misinformation-what-researchers-could-do-with-social-media-data/>, accessed 27 October 2023.

J.K. Paulus and D.M. Kent, 2020. “Predictably unequal: Understanding and addressing concerns that algorithmic clinical prediction may increase health disparities,” *NPJ Digital Medicine*, volume 3, number 1 (30 July), article number 99. doi: <https://doi.org/10.1038/s41746-020-0304-9>, accessed 27 October 2023.

S. Pichai, 2018. “AI at Google: Our principles” (7 June), at <https://blog.google/technology/ai/ai-principles/>, accessed 27 October 2023.

T. Pires, E. Schlinger, and D. Garrette, 2019. “How multilingual is multilingual BERT?” *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4,996–5,001, and at <https://aclanthology.org/P19-1493/>, accessed 27 October 2023.

N. Pourdamghani and K. Knight, 2019. “Neighbors helping the poor: Improving low-resource machine translation using related languages,” *Machine Translation*, volume 33, number 3, pp. 239–258. doi: <https://doi.org/10.1007/s10590-019-09236-7>, accessed 27 October 2023.

M.O.R. Prates, P.H. Avelar, and L.C. Lamb, 2020. “Assessing gender bias in machine translation: A case study with Google Translate,” *Neural Computing and Applications*, volume 32, pp. 6,363–6,381. doi: <https://doi.org/10.1007/s00521-019-04144-6>, accessed 27 October 2023.

A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, 2018. “Improving language understanding by generative pre-training,” at https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf, accessed 27 October 2023.

A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, 2019. “Language models are unsupervised multitask learners,” at <https://d4mucfpksyv.cloudfront.net/better-language->

[models/language_models_are_unsupervised_multitask_learners.pdf](#), accessed 27 October 2023.

C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P.J. Liu, 2020. “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal of Machine Learning Research*, volume 21, number 1, article number 140, pp. 5,485–5,551.

M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, 2020. “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 469–481.
doi: <https://doi.org/10.1145/3351095.3372828>, accessed 27 October 2023.

I.D. Raji, A. Smart, R.N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, and P. Barnes, 2020. “Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing,” *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 33–44.
doi: <https://doi.org/10.1145/3351095.3372873>, accessed 27 October 2023.

S. Ranathunga, E.-S. A. Lee, M. Prifti Skenduli, R. Shekhar, M. Alam, and R. Kaur, 2023. “Neural machine translation for low-resource languages: A survey,” *ACM Computing Surveys*, volume 55, number 11, article number 229, pp. 1–37.
doi: <https://doi.org/10.1145/3567592>, accessed 27 October 2023.

D. Reisman, J. Schultz, K. Crawford, and M. Whittaker, 2018. “Algorithmic impact assessments: A practical framework for public agency,” *AI Now* (9 April), at <https://ainowinstitute.org/publication/algorithmic-impact-assessments-report-2#>, accessed 27 October 2023.

M.T. Ribeiro, T. Wu, C. Guestrin, and S. Singh, 2020. “Beyond accuracy: Behavioral testing of NLP models with checklist,” *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4,902–4,912, and at <https://aclanthology.org/2020.acl-main.442/>, accessed 27 October 2023.

S.T. Roberts, 2019. *Behind the screen: Content moderation in the shadows of social media*. New Haven, Conn.: Yale University Press.

S. Ruder, I. Vulić, and A. Søgaard, 2019. “A survey of cross-lingual word embedding models,” *Journal of Artificial Intelligence Research*, volume 65, pp. 569–631.
doi: <https://doi.org/10.1613/jair.1.11640>, accessed 27 October 2023.

M. Sap, D. Card, S. Gabriel, Y. Choi, and N.A. Smith, 2019. “The risk of racial bias in hate speech detection,” *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1,668–1,678.
doi: <https://doi.org/10.18653/v1/P19-1163>, accessed 27 October 2023.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, 2017. “Proximal policy optimization algorithms,” *arXiv:1707.06347* (20 July).
doi: <https://doi.org/10.48550/arXiv.1707.06347>, accessed 27 October 2023.

R. Schwartz, J. Dodge, N.A. Smith, and O. Etzioni, 2020. “Green AI,” *Communications of the ACM*, volume 63, number 12, pp. 54–63.
doi: <https://doi.org/10.1145/3381831>, accessed 27 October 2023.

N. Selwyn, 2019. *Should robots replace teachers?: AI and the future of education*. Cambridge: Polity Press.

K. Shahriari and M. Shahriari, 2017. “IEEE standard review — ethically aligned design: A vision for prioritizing human well-being with artificial intelligence and autonomous systems,” *2017 IEEE Canada*

International Humanitarian Technology Conference (IHTC), pp. 197–201.
doi: <https://doi.org/10.1109/IHTC.2017.8058187>, accessed 27 October 2023.

H. Smith, 2021. “Clinical AI: Opacity, accountability, responsibility and liability,” *AI & Society*, volume 36, number 2, pp. 535–545.
doi: <https://doi.org/10.1007/s00146-020-01019-6>, accessed 27 October 2023.

N.A. Smith, 2020. “Contextual word representations: Putting words into computers,” *Communications of the ACM*, volume 63, number 6, pp. 66–74.
doi: <https://doi.org/10.1145/3347145>, accessed 27 October 2023.

N.A. Smuha, 2019. “The EU approach to ethics guidelines for trustworthy artificial intelligence,” *Computer Law Review International*, volume 20, number 4, pp. 97–106.
doi: <https://doi.org/10.9785/cr-2019-200402>, accessed 27 October 2023.

I. Solaiman, M. Brundage, J. Clark, A. Askeel, A. Herbert-Voss, J. Wu, A. Radford, G. Krueger, J.W. Kim, S. Kreps, M. McCain, A. Newhouse, J. Blazakis, K. McGuffie, and J. Wang, 2019. “Release strategies and the social impacts of language models,” *arXiv:1908.09203* (24 August).
doi: <https://doi.org/10.48550/arXiv.1908.09203>, accessed 27 October 2023.

A. Srivastava, A. Rastogi, A., Rao, A.A.M. Shoeb, A. Abid, A. Fisch, A.R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, A. Kluska, A. Lewkowycz, A. Agarwal, A. Power, A. Ray, A. Warstadt, A.W. Kocurek, A. Safaya, A. Tazarv, A. Xiang, A. Parrish, A. Nie, A. Hussain, A. Askeel, A. Dsouza, A. Slone, A. Rahane, A.S. Iyer, A. Andreassen, A. Madotto, A. Santilli, A. Stuhlmüller, A. Dai, A. La, A. Lampinen, A. Zou, A. Jiang, A. Chen, A. Vuong, A. Gupta, A. Gottardi, A. Norelli, A. Venkatesh, A. Gholamidavoodi, A. Tabassum, A. Menezes, A. Kirubarajan, A. Mullokandov, A. Sabharwal, A. Herrick, A. Efrat, A. Erdem, A. Karakaş, B.R. Roberts, B.S. Loe, B. Zoph, B. Bojanowski, B. Özyurt, B. Hedayatnia, B. Neyshabur, B. Inden, B. Stein, B. Ekmekci, B.Y. Lin, B. Howald, B. Orinion, C. Diao, C. Dour, C. Stinson, C. Argueta, C.F. Ramírez, C. Singh, C. Rathkopf, C. Meng, C. Baral, C. Wu, C. Callison-Burch, C. Waites, C. Voigt, C.D. Manning, C. Potts, C. Ramirez, C.E. Rivera, C. Siro, C. Raffel, C. Ashcraft, C. Garbacea, D. Sileo, D. Garrette, D. Hendrycks, D. Kilman, D. Roth, D. Freeman, D. Khashabi, D. Levy, D.M. González, D. Perszyk, D. Hernandez, D. Chen, D. Ippolito, D. Gilboa, D. Dohan, D. Drakard, D. Jurgens, D. Datta, D. Ganguli, D. Emelin, D. Kleyko, D. Yuret, D. Chen, D. Tam, D. Hupkes, D. Misra, D. Buzan, D.C. Mollo, D. Yang, D.-H. Lee, D. Schrader, E. Shutova, E.D. Cubuk, E. Segal, E. Hagerman, E. Barnes, E. Donoway, E. Pavlick, E. Rodola, E. Lam, E. Chu, E. Tang, E. Erdem, E. Chang, E.A. Chi, E. Dyer, E. Jerzak, E. Kim, E.E. Manyasi, E. Zheltonozhskii, F. Xia, F. Siar, F. Martínez-Plumed, F. Happé, F. Chollet, F. Rong, G. Mishra, G.I. Winata, G. de Melo, G. Kruszewski, G. Parascandolo, G. Mariani, G. Wang, G. Jaimovitch-López, G. Betz, G. Gur-Ari, H. Galijasevic, H. Kim, H. Rashkin, H. Hajishirzi, H. Mehta, H. Bogar, H. Shevlin, H. Schütze, H. Yakura, H. Zhang, H.M. Wong, I. Ng, I. Noble, J. Jumelet, J. Geissinger, J. Kernion, J. Hilton, J. Lee, J. Fernández Fisac, J.B. Simon, J. Koppel, J. Zheng, J. Zou, J. Kocoń, J. Thompson, J. Wingfield, J. Kaplan, J. Radom, J. Sohl-Dickstein, J. Phang, J. Wei, J. Yosinski, J. Novikova, J. Bosscher, J. Marsh, J. Kim, J. Taal, J. Engel, J. Alabi, J. Xu, J. Song, J. Tang, J. Waweru, J. Burden, J. Miller, J.U. Balis, J. Batchelder, J. Berant, J. Frohberg, J. Rozen, J. Hernandez-Orallo, J. Boudeman, J. Guerr, J. Jones, J.B. Tenenbaum, J.S. Rule, J. Chua, K. Kanclerz, K. Livescu, K. Krauth, K. Gopalakrishnan, K. Ignatyeva, K. Markert, K.D. Dhole, K. Gimpel, K. Omondi, K. Mathewson, K. Chiafullo, K. Shkaruta, K. Shridhar, K. McDonell, K. Richardson, L. Reynolds, L. Gao, L. Zhang, L. Dugan, L. Qin, L. Contreras-Ochando, L.-P. Morency, L. Moschella, L. Lam, L. Noble, L. Schmidt, L. He, L.O. Colón, L. Metz, L.K. Şenel, M. Bosma, M. Sap, M. ter Hoeve, M. Farooqi, M. Faruqui, M. Mazeika, M. Baturan, M. Marelli, M. Maru, M.J.R. Quintana, M. Tolkiehn, M. Giulianelli, M. Lewis, M. Potthast, M.L. Leavitt, M. Hagen, M. Schubert, M.O. Baitemirova, M. Arnaud, M. McElrath, M.A. Yee, M. Cohen, M. Gu, M. Ivanitskiy, M. Starritt, M. Strube, M. Swędrowski, M. Bevilacqua, M. Yasunaga, M. Kale, M. Cain, M. Xu, M. Suzgun, M. Walker, M. Tiwari, M. Bansal, M. Aminnaseri, M. Geva, M. Gheini, M. Varma T, N. Peng, N.A. Chi, N. Lee, N.G.-A. Krakover, N. Cameron, N. Roberts, N. Doiron, N. Martinez, N. Nangia, N. Deckers, N. Muennighoff, N.S. Keskar, N.S. Iyer, N. Constant, N. Fiedel, N. Wen, O. Zhang,

O. Agha, O. Elbaghdadi, O. Levy, O. Evans, P.A.M. Casares, P. Doshi, P. Fung, P.P. Liang, P. Vicol, P. Alipoormolabashi, P. Liao, P. Liang, P. Chang, P. Eckersley, P.M. Htut, P. Hwang, P. Miłkowski, P. Patil, P. Pezeshkpour, P. Oli, Q. Mei, Q. Lyu, Q. Chen, R. Banjade, R.E. Rudolph, R. Gabriel, R. Habacker, R. Risco, R. Millièrè, R. Garg, R. Barnes, R.A. Saurous, R. Arakawa, R. Raymaekers, R. Frank, R. Sikand, R. Novak, R. Sitelew, R. LeBras, R. Liu, R. Jacobs, R. Zhang, R. Salakhutdinov, R. Chi, R. Lee, R. Stovall, R. Teehan, R. Yang, S. Singh, S.M. Mohammad, S. Anand, S. Dillavou, S. Shleifer, S. Wiseman, S. Gruetter, S.R. Bowman, S.S. Schoenholz, S. Han, S. Kwatra, S.A. Rous, S. Ghazarian, S. Ghosh, S. Casey, S. Bischoff, S. Gehrmann, S. Schuster, S. Sadeghi, S. Hamdan, S. Zhou, S. Srivastava, S. Shi, S. Singh, S. Asaadi, S.S. Gu, S. Pachchigar, S. Toshniwal, S. Upadhyay, S.(S.) Debnath, S. Shakeri, S. Thormeyer, S. Melzi, S. Reddy, S.P. Makini, S.-H. Lee, S. Torene, S. Hatwar, S. Dehaene, S. Divic, S. Ermon, S. Biderman, S. Lin, S. Prasad, S.T. Piantadosi, S.M. Shieber, S. Mishnerghi, S. Kiritchenko, S. Mishra, T. Linzen, T. Schuster, T. Li, T. Yu, T. Ali, T. Hashimoto, T.-L. Wu, T. Desbordes, T. Rothschild, T. Phan, T. Wang, T. Nkinyili, T. Schick, T. Kornev, T. Tunduny, T. Gerstenberg, T. Chang, T. Neeraj, T. Khot, T. Shultz, U. Shaham, V. Misra, V. Demberg, V. Nyamai, V. Raunak, V. Ramasesh, V.U. Prabhu, V. Padmakumar, V. Srikumar, W. Fedus, W. Saunders, W. Zhang, W. Vossen, X. Ren, X. Tong, X. Zhao, X. Wu, X. Shen, Y. Yaghoobzadeh, Y. Lakretz, Y. Song, Y. Bahri, Y. Choi, Y. Yang, Y. Hao, Y. Chen, Y. Belinkov, Y. Hou, Y. Bai, Z. Seid, Z. Zhao, Z. Wang, Z.J. Wang, Z. Wang, and Z. Wu, 2022. “Beyond the imitation game: Quantifying and extrapolating the capabilities of language models,” *arXiv:2206.04615* (9 June).

doi: <https://doi.org/10.48550/arXiv.2206.04615>, accessed 27 October 2023.

J. Stoyanovich, B. Howe, and H. Jagadish, 2020. “Responsible data management,” *Proceedings of the VLDB Endowment*, volume 13, number 12, pp. 3474–3.488.

doi: <https://doi.org/10.14778/3415478.3415570>, accessed 27 October 2023.

T. Sun, A. Gaut, S. Tang, Y. Huang, M. ElSherief, J. Zhao, D. Mirza, E. Belding, K.-W. Chang, and W.Y. Wang, 2019. “Mitigating gender bias in natural language processing: Literature review,” *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1,630–1,640, and at <https://aclanthology.org/P19-1159/>, accessed 27 October 2023.

I. Sutskever, O. Vinyals, and Q.V. Le, 2014. “Sequence to sequence learning with neural networks,” *NIPS’14: Proceedings of the 27th International Conference on Neural Information Processing Systems*, volume 2, pp. 3,104–3,112, and at https://papers.nips.cc/paper_files/paper/2014, accessed 27 October 2023.

M. Taddeo and L. Floridi, 2018. “Regulate artificial intelligence to avert cyber arms race,” *Nature*, volume 556, number 7701 (19 April), pp. 296–298.

doi: <https://doi.org/10.1038/d41586-018-04602-6>, accessed 27 October 2023.

L. Taylor and R. Schroeder, 2015. “Is bigger better? The emergence of big data as a tool for international development policy,” *GeoJournal*, volume 80, pp. 503–518.

doi: <https://doi.org/10.1007/s10708-014-9603-5>, accessed 27 October 2023.

A. Torralba and A.A. Efros, 2011. “Unbiased look at dataset bias,” *CVPR 2011: Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1,521–1,528.

doi: <https://doi.org/10.1109/CVPR.2011.5995347>, accessed 27 October 2023.

H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample, 2023. “LLaMA: Open and efficient foundation language models,” *arXiv:2302.13971* (27 February).

doi: <https://doi.org/10.48550/arXiv.2302.13971>, accessed 27 October 2023.

H.C. Triandis, 1995. *Individualism and collectivism*. New York: Routledge.

doi: <https://doi.org/10.4324/9780429499845>, accessed 27 October 2023.

B. Ustun, A. Spangher, and Y. Liu, 2019. “Actionable recourse in linear classification,” *FAT* ’19*:

Proceedings of the Conference on Fairness, Accountability, and Transparency, pp. 10–19.
doi: <https://doi.org/10.1145/3287560.3287566>, accessed 27 October 2023.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, and I. Polosukhin, 2017. “Attention is all you need,” *Advances in Neural Information Processing Systems*, volume 30, at https://papers.nips.cc/paper_files/paper/2017, accessed 27 October 2023.

S. Wachter, B. Mittelstadt, and L. Floridi, 2017. “Transparent, explainable, and accountable AI for robotics,” *Science Robotics*, volume 2, number 6 (31 May).
doi: <https://doi.org/10.1126/scirobotics.aan60>, accessed 27 October 2023.

A. Wallach, C. Allen, and I. Smit, 2008. “Machine morality: Bottom-up and top-down approaches for modelling human moral faculties,” *AI & Society*, volume 22, pp. 565–582.
doi: <https://doi.org/10.1007/s00146-007-0099-0>, accessed 27 October 2023.

Q. Wang, B. Li, T. Xiao, J. Zhu, C. Li, D.F. Wong, and L.S. Chao, 2019. “Learning deep transformer models for machine translation,” *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1,810–1,822, and at <https://aclanthology.org/P19-1176.pdf>, accessed 27 October 2023.

J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E.H. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, and W. Fedus 2022. “Emergent abilities of large language models,” *arXiv:2206.07682* (15 June).
doi: <https://doi.org/10.48550/arXiv.2206.07682>, accessed 27 October 2023.

J. Wei, X. Wang, D. Schuurmans, M. Bosma, E. Chi, Q. Le, and D. Zhou, 2022. “Chain of thought prompting elicits reasoning in large language models,” *arXiv:2201.11903* (28 January).
doi: <https://doi.org/10.48550/arXiv.2201.11903>, accessed 27 October 2023.

L. Weidinger, J. Mellor, M. Rauh, C. Griffin, J. Uesato, P.-S. Huang, M. Cheng, M. Glaese, B. Balle, A. Kasirzadeh, Z. Kenton, S. Brown, W. Hawkins, T. Stepleton, C. Biles, A. Birhane, J. Haas, L. Rimell, L.A. Hendricks, W. Isaac, S. Legassick, G. Irving, and I. Gabriel, 2021. “Ethical and social risks of harm from language models,” *arXiv:2112.04359* (8 December).
doi: <https://doi.org/10.48550/arXiv.2112.04359>, accessed 27 October 2023.

T.-H. Wen, D. Vandyke, N. Mrkšić, M. Gašić, L.M. Rojas-Barahona, P.-H. Su, S. Ultes, and S. Young, 2017. “A network-based end-to-end trainable task-oriented dialogue system,” *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 1, Long papers, pp. 438–449, and at <https://aclanthology.org/E17-1042/>, accessed 27 October 2023.

S.M. West, M. Whittaker, and K. Crawford, 2019. “Discriminating systems: Gender, race, and power in AI,” *AI Now* (1 April), at <https://ainowinstitute.org/publication/discriminating-systems-gender-race-and-power-in-ai-2>, accessed 27 October 2023.

B.L. Whorf, 1964. *Language, thought, and reality: Selected Writings of Benjamin Lee Whorf*. Edited by John B. Carroll. Cambridge, Mass.: MIT Press.

C. Wu, S. Yin, W. Qi, X. Wang, Z. Tang, and N. Duan, 2023. “Visual ChatGPT: Talking, drawing and editing with visual foundation models,” *arXiv:2303.04671* (8 March).
doi: <https://doi.org/10.48550/arXiv.2303.04671>, accessed 27 October 2023.

S. Yan, H.-t. Kao, and E. Ferrara, 2020. “Fair class balancing: Enhancing model fairness without observing sensitive attributes,” *CIKM '20: Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pp. 1,715–1,724.
doi: <https://doi.org/10.1145/3340531.3411980>, accessed 27 October 2023.

K.-C. Yang, E. Ferrara, and F. Menczer, 2022. “Botometer 101: Social bot practicum for computational social scientists,” *Journal of Computational Social Science*, volume 5, number 2, pp. 1,511–1,528. doi: <https://doi.org/10.1007/s42001-022-00177-5>, accessed 27 October 2023.

K. Yeung, 2020. “Recommendation of the council on artificial intelligence (OECD),” *International Legal Materials*, volume 59, number 1, pp. 27–34. doi: <https://doi.org/10.1017/ilm.2020.5>, accessed 27 October 2023.

T. Young, D. Hazarika, S. Poria, and E. Cambria, 2018. “Recent trends in deep learning based natural language processing,” *IEEE Computational Intelligence*, volume 13, number 3, pp. 55–75. doi: <https://doi.org/10.1109/MCI.2018.2840738>, accessed 27 October 2023.

M.B. Zafar, I. Valera, M. Gomez Rodriguez, and K.P. Gummadi, 2017. “Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment,” *WWW '17: Proceedings of the 26th International Conference on World Wide Web*, pp. 1,171–1,180. doi: <https://doi.org/10.1145/3038912.3052660>, accessed 27 October 2023.

R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi, 2019. “Defending against neural fake news,” *ANIPS'19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, article number 812, pp. 9,054–9,065.

N.H. Zhang, B. Lemoine, and M. Mitchell, 2018. “Mitigating unwanted biases with adversarial learning,” *AIES '18: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 335–340. doi: <https://doi.org/10.1145/3278721.3278779>, accessed 27 October 2023.

S. Zhang, E. Dinan, J. Urbanek, A. Szlam, D. Kiela, and J. Weston, 2018. “Personalizing dialogue agents: I have a dog, do you have pets too?” *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, volume 1: Long papers, pp. 2,204–2,213, and at <https://aclanthology.org/P18-1205/>, accessed 27 October 2023.

D.M. Ziegler, N. Stiennon, J. Wu, T.B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving, 2019. “Fine-tuning language models from human preferences,” *arXiv:1909.08593* (18 September). doi: <https://doi.org/10.48550/arXiv.1909.08593>, accessed 27 October 2023.

Editorial history

Received 14 October 2023; revised 23 October 2023; accepted 25 October 2023.



This paper is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

Should ChatGPT be biased? Challenges and risks of bias in large language models
by Emilio Ferrara.

First Monday, volume 28, number 11 (November 2023).

doi: <https://dx.doi.org/10.5210/fm.v28i11.13346>