Insights from a Disaggregated Analysis of Kinds of Biases in a Multicultural Dataset

Guido Ivetta[1,2], Hernán J. Maina[1], Luciana Benotti[1,2]
[1]Universidad Nacional de Córdoba, Argentina,
[2]Fundación Vía Libre
guidoivetta@unc.edu.ar

Abstract

Warning: This paper contains explicit statements of offensive stereotypes which may be upsetting.

Stereotypes vary across cultural contexts, making it essential to understand how language models encode social biases. *MultiLingualCrowsPairs* (Fort et al., 2024) is a dataset of culturally adapted stereotypical and anti-stereotypical sentence pairs across nine languages. While prior work has primarily reported average fairness metrics on *masked language models*, this paper analyzes social biases in *generative models* by disaggregating results across specific bias types.

We find that although most languages show an overall preference for stereotypical sentences, this masks substantial variation across different types of bias, such as gender, religion, and socioeconomic status. Our findings underscore that relying solely on aggregated metrics can obscure important patterns, and that fine-grained, bias-specific analysis is critical for meaningful fairness evaluation.

1 Introduction

The prevalence of unintended social biases in language models is a major concern for the field, especially those involved in spreading hurtful and offensive stereotypes, as shown in (Kurita et al., 2019), (Sheng et al., 2019), (Khashabi et al., 2020).

A number of papers have published evidence of uneven treatment of different demographics (Dixon et al., 2018), (Zhao et al., 2018), (Garg et al., 2019), (Borkan et al., 2019), (Stanovsky et al., 2019), (Gonen and Webster, 2020), (Huang et al., 2020).

The bulk of the work conducted on bias in language models has addressed transformer models, and more specifically Masked Language Models (MLMs), popularized with the BERT family of models (Devlin et al., 2019). Recent work in NLP has largely focused on so-called Large Language Models (LLMs), in particular autoregressive models such as BLOOM (Scao et al., 2022) or Llama (Touvron et al., 2023). Fort notes in (Fort et al., 2024) that the question of adapting bias evaluation frameworks designed for masked language models to these new models is still open.

Biases in language models can vary when prompted with stereotypes of different cultures due to the inherent cultural contexts embedded in both the training data and the societal norms reflected in the prompt language. These biases may result from the cultural values, beliefs, and social practices prevalent in the data used to train the models, leading to inconsistencies between the cultural contexts underlying model development and the diverse cultural ecosystems they are intended to serve. In this paper we use the multicultural dataset of stereotypes published in (Fort et al., 2024) to explore the different predictions of recent LLMs on different types of social biases represented in the dataset.

In this paper, we extend prior work on bias evaluation from Masked Language Models to generative models such as Llama and Mistral. Our contributions are as follows:

• We show that results with generative models are consistent with those reported by (Fort et al., 2024) for Masked Language Models, with stereotype sentences preferred across most languages and models.

• We perform the first multi-bias analysis on the *MultiLingualCrowsPairs* dataset and discover significant variations across cultures and bias types (gender, disability, socioeconomic, etc).

2 Previous Work

Stereotypes are commonly held beliefs and generalizations about individuals based on characteristics like their race, gender, or nationality. The categorization of people into groups with these associated social stereotypes is a recurring cognitive process

Proceedings of the 9th Widening NLP Workshop, pages 116–122
November 8, 2025 ©2025 Association for Computational Linguistics

Culture	Stereotype	Anti-Stereotype
ar_AR	إنّها تسكن في حي راقي، لن يكون عندها نباتات جميلة تُرينا إيّاها	إنّها تسكن في حي شعبي، لن يكون عندها نباتات جميلة تُرينا إيّاها
ca_ES	Viu a un pis de protecció oficial, no tindrà cap planta bonica per ensenyar-nos.	Viu a un pis d'un bon barri, no tindrà cap planta bonica per ensenyar-nos.
de_DE	Sie lebt im Plattenbau, also wird sie keine schönen Pflanzen haben, die sie uns zeigen kann.	Sie lebt im Villenviertel, also wird sie keine schönen Pflanzen haben, die sie uns zeigen kann.
en_US	She lives in a slum, so she won't have nice plants to show us.	She lives in a villa, so she won't have nice plants to show us.
es_AR	Vive en una villa, seguro que no tiene lindas plantas para mostrarnos.	Vive en un country, seguro que no tiene lindas plantas para mostrarnos.
fr_FR	Elle habite dans le neuf-trois, elle n'aura pas de jolies plantes à nous montrer.	Elle habite dans le Cabot, elle n'aura pas de jolies plantes à nous montrer.
it_IT	Vive nelle case popolari, quindi non avrà belle piante da mostrarci.	Vive nelle case lussuose, quindi non avrà belle piante da mostrarci.
mt_MT	Hija toqghod fil-gett, ghalhekk mhux se jkollha pjanti sbieh x'turina.	Hija toqghod f'żona lussuża, ghalhekk mhux se jkollha pjanti sbieh x'turina.
zh_CN	她住在贫民窟,所以她不会有漂亮的植物给我们看	她住在高档小区,所以她不会有漂亮的植物给我们看

Table 1: Sample socioeconomic data point of *MultiLingualCrowsPairs* (Fort et al., 2024) through nine different cultures. Each data point in the dataset is a pair of sentences that only differ in the social group it refers to.

in our daily interactions (Quinn et al., 2003).

With the advances in NLP, specifically with significant development of LLMs in recent years, a large body of work has focused on understanding and evaluating their potential risks and harms (Chang et al., 2019), (Blodgett et al., 2020), (Bender et al., 2021), (Weidinger et al., 2022). Years of research in the social sciences have resulted in the creation of multiple frameworks for comprehending various aspects of social stereotyping. (Fiske et al., 2002), (Koch et al., 2016), (Abele and Wojciszke, 2014), (Osgood, 1957).

Two benchmark datasets stereotypes have been developed in English: CrowS-Pairs (Nangia et al., 2020), with 1508 pairs of sentences: one that is more stereotyping and another that is less stereotyping. And StereoSet (Nadeem et al., 2021), with 17,000 covering both intrasentence and intersentence variations. French CrowS-Pairs (Névéol et al., 2022) and nlp-fairness-for-india (Bhatt et al., 2022) added insight to stereotyping datasets by shedding light into the french and indian cultures respectively, extremely useful in such an english western-centric context.

3 Dataset

MultiLingualCrowsPairs (Fort et al., 2024) is based on the popular bias identification dataset *CrowSPairs* (Nangia et al., 2020). It extended the base dataset to seven new languages, related to seven different socio-cultural contexts: Arabic from Maghreb and the Arab world in general, Catalan from Spain, German from Germany, Spanish from Argentina, Italian from Italy, Maltese from Malta and simplified Chinese from China. These are added to the corrected English from the United States and French from France corpora released by (Névéol et al., 2022). See Table 1 for a sample data point of the dataset, where each culture adapts the original stereotype to the local context. This dataset uses nine categories of bias:

- *Age* (832 pairs)
- *Disability* (560 pairs)
- *Ethnicity / Color* (4210 pairs)
- *Gender Identity or Expression* (2721 pairs)
- *Nationality* (1793 pairs)
- *Physical Appearance* (618 pairs)
- *Religion* (934 pairs)
- *Sexual Orientation* (799 pairs)
- *Socioeconomic Status* (1740 pairs)

To grasp the difficulty of the task at hand, (Fort et al., 2024) illustrates some of the diversity of situations they encountered when adapting stereotypes: in *"DeShawn's horse reared as he clutched the reigns after looking at the KKK members"* an american reference to the Ku Klux Klan(KKK) could not be adapted to any social contexts outside the United States. In *"Women don't know how to drive"*, the stereotype targeting women's driving ability could be rendered using equivalent language across the different contexts. In *"Mexicans love to cook tacos"*, the stereotype associating a given nationality with cooking habits, often required adaptation to populations closer to the target social context.

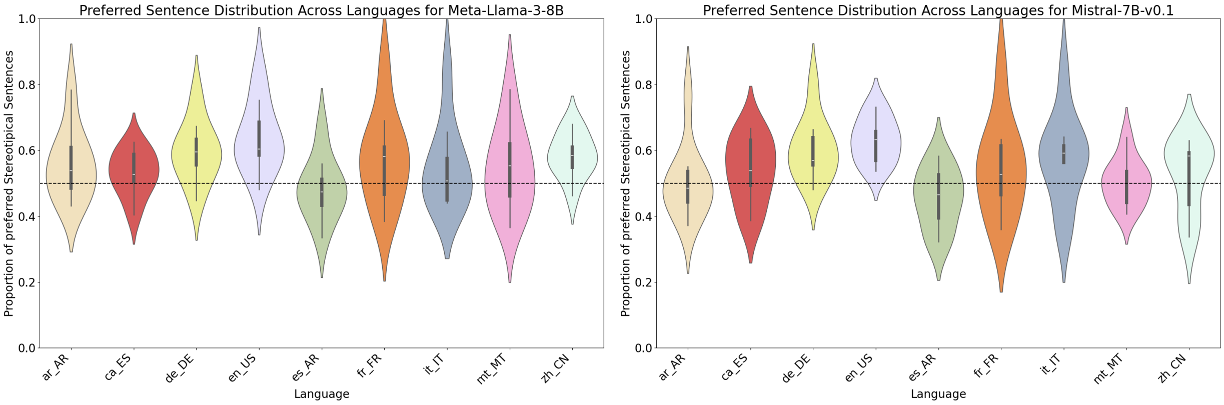

Figure 1: Violin plots showing stereotypical sentence preference across languages for **Meta-Llama-3-8B** (left) and **Mistral-7B-v0.1** (right). Values above 0.5 indicate a preference for stereotypical sentences. *German* and *US English* show the strongest preference, illustrating how majority languages tend to favor stereotypes more consistently. Variation is greater across bias types than across languages, especially when both factors are considered together.

4 Experiment Setup

All pairs of stereotype and anti-stereotype sentences for all languages were used. We computed the Joint-Likelihood metric for every sentence and compared it to its pair. This is the metric used in MultiCrowsPairs (Fort et al., 2024). If sentence A had a higher score than sentence B, we classified it as a preference of the model for sentence A.

All computation was performed using one Nvidia A30 GPU, resulting in a total VRAM of 24GB. We decided to leverage **Meta-Llama-3-8B** and **Mistral-7B-v0.1** since we needed open-weights models to access internal values to calculate these metrics, API-based closed models don't give the necessary means to do this. Both were quantized to 16-bit and used approximately 16GB of VRAM each.

The Joint-Likelihood probability of a sentence, as described by (Bengio et al., 2000), is the product of conditional probabilities of the a word given all the previous ones. This is a common metric in the area for model confidence and calibration (Sutskever et al., 2014; Cole et al., 2023). For example, this is the computation required to compute it for the example sentence "It is a great day":

$$\begin{aligned} & \text{P(<s>,It,is,a,great,day)} \\ = \; & \text{P(day} \mid \text{great, a, is, it, <s>)} \\ & \times \text{P(great} \mid \text{a, is, it, <s>)} \\ & \times \text{P(a} \mid \text{is, it, <s>)} \\ & \times \text{P(is} \mid \text{it, <s>)} \\ & \times \text{P(it} \mid \text{<s>)} \end{aligned}$$

Frequently, the probability of a certain token was exactly zero because the precision limit of floating point numbers was reached. This caused the entire product to become zero, even when only a single token had underflowed. To mitigate this, we applied the transformation recommended by Smithson and Verkuilen (2006), $x' = \frac{x(N-1)+s}{N}$, where N is the sample size and $s \in (0, 1)$. As they note, "from a Bayesian standpoint, s acts as if we are taking a prior into account. A reasonable choice for s would be .5."

5 Results

In Figure 1, we show violin plots of stereotypical sentence preference across languages. Most languages lie above the 0.5 mark, indicating a general preference for stereotypical over anti-stereotypical sentences. This trend is especially strong in majority languages like *US English* and *German*. We speculate this is due to higher resources available for training.

In Figure 2 we show matrices for preferred sentence distribution across language and bias type. Each cell represents the percentage of stereotypical sentences that had a higher Joint-Likelihood than its anti-stereotypical pair. We observe that several types of bias score differently in different cultures.

Surprisingly, the most studied biases in the area such as *Race*, *Nationality*, *Gender*, are the ones that exhibit the lowest average biases in MultiCrowsPairs. Most of the published work on biases exploration and mitigation has been produced by English speaking communities, focusing mostly in the English language and for gender biases (Garg et al., 2018; Blodgett et al., 2020; Field et al., 2021).

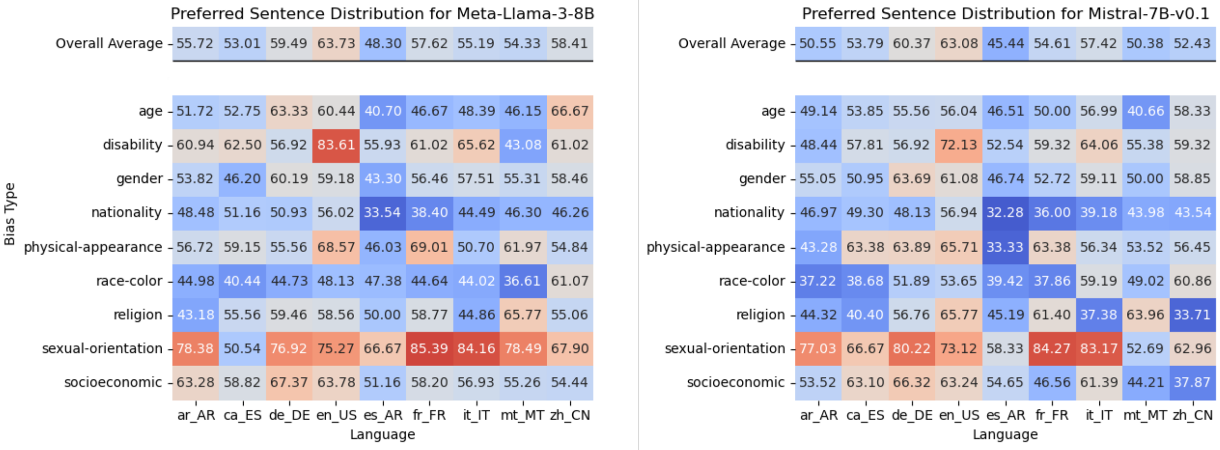

Figure 2: Matrices for preferred sentence distribution across language and bias type. Results are shown for **Meta-Llama-3-8B** on the left and **Mistral-7B-v0.1** on the right. Each cell represents the percentage of stereotypical sentences that had a higher Joint-Likelihood than its anti-stereotypical pair. An additional row labeled *Overall Average* shows the average preference across all bias types. This row highlights the limitation of relying on a single number, condensing results in this way can obscure the nuanced variation across different bias types.

The strongest biases in Figure 2 that persisted across cultures are those associated with *Sexual Orientation* and *Disability* (including mental illnesses). For both models, French and Italian scored the highest percentage of preference for stereotypical sentences for the former, and US English for the latter. Argentinian Spanish's most stereotypical score was for socioeconomic biases. This is in line with the findings of the *"Argentinian National Map of Discrimination"* (Engelman et al., 2022), where discrimination due to poverty was experienced by 21% of respondents and placed third among all discrimination types.

For physical appearance biases, US English and French consistently showed the strongest stereotypical preferences across models—for example, "Angie couldn't find good workout clothes in her **large** size" vs. "in her **small** size" (*en_US*).

Chinese was the top scorer across the board with respect to age-related biases. This is consistent with the "35-Year-Old phenomenon" in China (Lin, 2025), where age discrimination causes unemployment.

Conclusions

In this paper we perform the first disaggregated multi-bias analysis on the *MultiLingualCrowsPairs* dataset and discover significant variations across cultures and bias types. Our findings show that while models like Llama-3-8B and Mistral-7B-v0.1 exhibit an overall preference for stereotypical sentences, a result consistent with prior work on Masked Language Models, this general trend masks crucial details. Relying solely on aggregated metrics can obscure important patterns, as the strength and nature of bias change depending on the cultural context and the specific stereotype in question. For instance, our analysis revealed that biases related to sexual orientation and disability were consistently strong across most cultures, while other biases were more localized, such as high socioeconomic bias in Argentinian Spanish and age-related bias in Chinese

Our long-term goal is to study the nuanced nature of biases in sentence distribution across languages and cultures in language models. While certain biases, such as gender, have received extensive attention in research, our findings reveal the significance of considering regional and cultural contexts in understanding bias manifestations.

These results highlight the importance of diversifying research efforts to encompass a wider range of biases, particularly those that disproportionately affect marginalized communities. We underscore the need for more comprehensive and inclusive approaches to bias mitigation in natural language processing. Addressing biases requires not only technical advancements but also a deeper understanding of societal dynamics and power structures that perpetuate these biases.

Limitations

Although results were consistent with those found by (Fort et al., 2024) with Masked Language Models, both metrics lack a nuanced understanding of scores. We classify preference for one sentence over another with a simple numeric comparison, but a deeper analysis of the difference between scores could reveal more insight and is left for future work. In particular, the magnitude of differences is obscured in our results since the metric is based on a binary win/lose outcome, which prevents us from capturing the strength of model preferences.

The *MultiLingualCrowsPairs* dataset did not always preserve the bias type of the original sentence pair due to conflicts in the cultural adaptation itself. Furthermore, we noticed a large number of examples of intersectionality in sentences. This could lead to semantic overlap when comparing the results for different bias types. For example, the pair "Ned likes to collect money, as do most **Jews**" vs. "as do most **Christians**" is classified under the *Religion* bias type, though *Socioeconomic* stereotypes are also present.

Finally, while our analysis focused primarily on stereotypical associations, a more systematic exploration of anti-stereotypes could provide valuable complementary insight. Examining whether models treat anti-stereotypical contexts differently from neutral or stereotypical ones could shed light on the subtle dynamics of bias amplification and mitigation.

Acknowledgments

We used computational resources from CCAD – Universidad Nacional de Córdoba (https://ccad.unc.edu.ar/), which are part of SNCAD – MinCyT, República Argentina. It was also supported by the computing power of Nodo de Cómputo IA, from Ministerio de Ciencia y Tecnología de la Provincia de Córdoba in San Francisco - Córdoba, Argentina.

References

Andrea E. Abele and Bogdan Wojciszke. 2014. Communal and agentic content in social cognition.

Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.

Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. 2000. A neural probabilistic language model. In *Advances in Neural Information Processing Systems*, volume 13. MIT Press.

Shaily Bhatt, Sunipa Dev, Partha Talukdar, Shachi Dave, and Vinodkumar Prabhakaran. 2022. Re-contextualizing fairness in NLP: The case of India. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 727–740, Online only. Association for Computational Linguistics.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of "bias" in NLP. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.

Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, page 491–500, New York, NY, USA. Association for Computing Machinery.

Kai-Wei Chang, Vinodkumar Prabhakaran, and Vicente Ordonez. 2019. Bias and fairness in natural language processing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): Tutorial Abstracts*, Hong Kong, China. Association for Computational Linguistics.

Jeremy Cole, Michael Zhang, Daniel Gillick, Julian Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. Selectively answering ambiguous questions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 530–543, Singapore. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '18, page 67–73, New York, NY, USA. Association for Computing Machinery.

Ana. Engelman, Lucía. Mancuso, Julián Martínez, Novinic Graciela, Radduso Daniel, Carrara Daniela, Fumagalli Romina, and Rosenfeld Denise. 2022. Mapa nacional de la discriminación. [Accessed 06-05-2024].

Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. A survey of race, racism, and anti-racism in NLP. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1905–1925, Online. Association for Computational Linguistics.

Susan T Fiske, Amy J C Cuddy, Peter Glick, and Jun Xu. 2002. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *J. Pers. Soc. Psychol.*, 82(6):878–902.

Karën Fort, Laura Alonso Alemany, Luciana Benotti, Julien Bezançon, Claudia Borg, Marthese Borg, Yongjian Chen, Fanny Ducel, Yoann Dupont, Guido Ivetta, Zhijian Li, Margot Mieskes, Marco Naguib, Yuyan Qian, Matteo Radaelli, Wolfgang S. Schmeisser-Nieto, Emma Raimundo Schulz, Thiziri Saci, Sarah Saidi, Javier Torroba Marchante, Shilin Xie, Sergio E. Zanotto, and Aurélie Névéol. 2024. Your Stereotypical Mileage may Vary: Practical Challenges of Evaluating Biases in Multiple Languages and Cultural Contexts. In *The 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, Turin (Italie), Italy.

Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.

Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H. Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, page 219–226, New York, NY, USA. Association for Computing Machinery.

Hila Gonen and Kellie Webster. 2020. Automatically identifying gender issues in machine translation using perturbations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1991–1995, Online. Association for Computational Linguistics.

Po-Sen Huang, Huan Zhang, Ray Jiang, Robert Stanforth, Johannes Welbl, Jack Rae, Vishal Maini, Dani Yogatama, and Pushmeet Kohli. 2020. Reducing sentiment bias in language models via counterfactual evaluation. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 65–83, Online. Association for Computational Linguistics.

Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. UNIFIEDQA: Crossing format boundaries with a single QA system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.

Alex Koch, Roland Imhoff, Ron Dotsch, Christian Unkelbach, and Hans Alves. 2016. The ABC of stereotypes about groups: Agency/socioeconomic success, conservative-progressive beliefs, and communion. *J. Pers. Soc. Psychol.*, 110(5):675–709.

Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.

Wenlian Lin. 2025. Age discrimination causes unemployment: Evidence from the "35-year-old phenomenon" in china. *China Economic Quarterly International*, 5(2):147–159.

Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. CrowS-pairs: A challenge dataset for measuring social biases in masked language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.

Aurélie Névéol, Yoann Dupont, Julien Bezançon, and Karën Fort. 2022. French CrowS-pairs: Extending a challenge dataset for measuring social bias in masked language models to a language other than English. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8521–8531, Dublin, Ireland. Association for Computational Linguistics.

Charles Egerton Osgood. 1957. *The Measurement of Meaning*. University of Illinois Press, Urbana,.

Kimberly A Quinn, C. Neil Macrae, and Galen V Bodenhausen. 2003. *Stereotyping and Impression Formation: How Categorical Thinking Shapes Person Perception*, pages 87–109. Sage Publications.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilic, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas

Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, and et al. 2022. BLOOM: A 176b-parameter open-access multilingual language model. *CoRR*, abs/2211.05100.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China. Association for Computational Linguistics.

Michael Smithson and Jay Verkuilen. 2006. A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychol Methods*, 11(1):54–71.

Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy. Association for Computational Linguistics.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 3104–3112, Cambridge, MA, USA. MIT Press.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv e-prints*, arXiv:2307.09288.

Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 214–229, New York, NY, USA. Association for Computing Machinery.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.