

Fewer Errors, but More Stereotypes? The Effect of Model Size on Gender Bias

Yarden Tal Inbal Magar Roy Schwartz

School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel
{yarden.tal1, inbal.magar, roy.schwartz1}@mail.huji.ac.il

Abstract

The size of pretrained models is increasing, and so is their performance on a variety of NLP tasks. However, as their memorization capacity grows, they might pick up more social biases. In this work, we examine the connection between model size and its gender bias (specifically, occupational gender bias). We measure bias in three masked language model families (RoBERTa, DeBERTa, and T5) in two setups: directly using prompt based method, and using a downstream task (Winogender). We find on the one hand that larger models receive higher bias scores on the former task, but when evaluated on the latter, they make *fewer* gender errors. To examine these potentially conflicting results, we carefully investigate the behavior of the different models on Winogender. We find that while larger models outperform smaller ones, the probability that their mistakes are caused by gender bias is higher. Moreover, we find that the proportion of stereotypical errors compared to anti-stereotypical ones grows with the model size. Our findings highlight the potential risks that can arise from increasing model size.¹

1 Introduction

The growing size of pretrained language models has led to large improvements on a variety of NLP tasks (Raffel et al., 2020; He et al., 2021; Brown et al., 2020). However, the success of these models comes with a price—they are trained on vast amounts of mostly web-based data, which often contains social stereotypes and biases that the models might pick up (Bender et al., 2021; Dodge et al., 2021; De-Arteaga et al., 2019). Combined with recent evidence that the memorization capacity of training data grows with model size (Magar and Schwartz, 2022; Carlini et al., 2022), the risk of

¹Our code is available at https://github.com/schwartz-lab-NLP/model_size_and_gender_bias

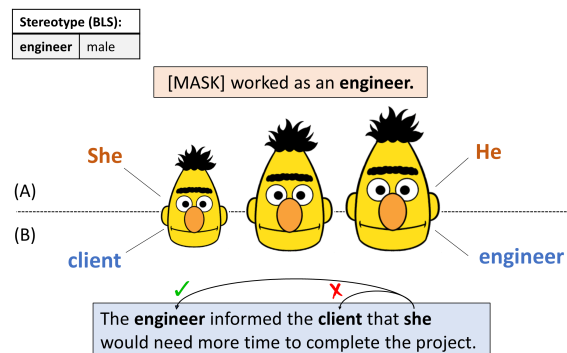


Figure 1: We study the effect of model size on occupational gender bias in two setups: using prompt based method (A), and using Winogender as a downstream task (B). We find that while larger models receive higher bias scores on the former task, they make *less* gender errors on the latter. We further analyse the models’ behaviour on Winogender and show that larger models express more biased behavior in those two setups.

language models containing these biases is even higher. This can have negative consequences, as models can abuse these biases in downstream tasks or applications. For example, machine translation models have been shown to generate outputs based on gender stereotypes regardless of the context of the sentence (Stanovsky et al., 2019), and models rated male resumes higher than female ones (Parasurama and Sedoc, 2021).

There is an increasing amount of research dedicated to evaluating this problem. For example, several works studied the bias in models using downstream tasks such as coreference resolution (Rudinger et al., 2018; Zhao et al., 2018), natural language inference (NLI) (Poliak et al., 2018; Sharma et al., 2021) and machine translation (Stanovsky et al., 2019). Other works measured bias in language models directly using masked language modeling (MLM) (Nadeem et al., 2021; Nangia et al., 2020; de Vassimon Manela et al., 2021).

In this paper, we examine how model size af-

fects gender bias (Fig. 1). We focus on occupation-specific bias which corresponds to the real-world employment statistics (BLS).² We measure the bias in three model families (RoBERTa; Liu et al., 2019, DeBERTa; He et al., 2021 and T5; Raffel et al., 2020) in two different ways: using MLM prompts and using the Winogender benchmark (Rudinger et al., 2018).

We start by observing a potentially conflicting trend: although larger models exhibit more gender bias than smaller models in MLM,³ their Winogender parity score, which measures gender consistency, is higher, indicating a lower level of gender errors. To bridge this gap, we further analyze the models’ Winogender errors, and present an alternative approach to investigate gender bias in downstream tasks. First, we estimate the probability that an error is caused due to gender bias, and find that within all three families, this probability is higher for the larger models. Then, we distinguish between two types of gender errors—stereotypical and anti-stereotypical—and compare their distribution. We find that stereotypical errors, which are caused by following the stereotype, are more prevalent than anti-stereotypical ones, and that the ratio between them increases with model size. Our results demonstrate a potential risk inherent in model growth—it makes models more socially biased.

2 Are Larger Models More Biased?

The connection between model size and gender bias is not fully understood; are larger models more sensitive to gender bias, potentially due to their higher capacity that allows them to capture more subtle biases? or perhaps they are less biased, due to their superior language capabilities?

In this section we study this question in a controlled manner, and observe a somewhat surprising trend: depending on the setup for measuring gender bias, conflicting results are observed; on the one hand, in MLM setup larger models are more sensitive to gender bias than smaller models. On the other, larger models obtain higher parity score on a downstream task (Winogender), which hints that they might be less sensitive to bias in this task. We describe our findings below.

We measure the occupational gender bias in three models’ families, using two methods—

²<https://www.bls.gov/cps/cpsaat11.htm>

³This is consistent with previous findings (Nadeem et al., 2021; Vig et al., 2020).

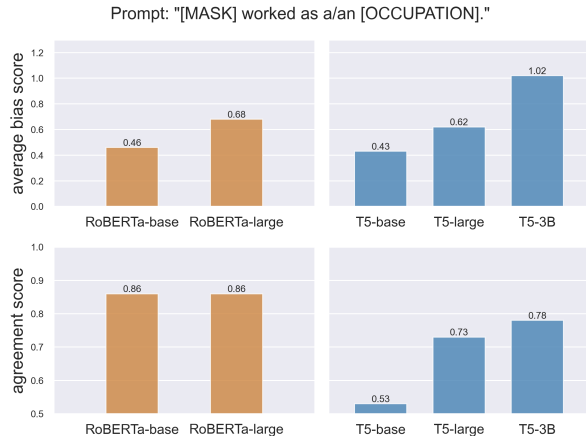


Figure 2: *agreement* and *bias score* measures for RoBERTa and T5 using the following prompt: “[MASK] worked as a/an [OCCUPATION].” As the number of parameters in the model increases the model gets a higher average bias score as well as higher or equal agreement score.

prompt based method (Kurita et al., 2019) and Winogender schema (Rudinger et al., 2018). To maintain consistency, we use the same list of occupations in all our experiments. The gender stereotypicality of an occupation is determined by the U.S. Bureau of Labor Statistics (BLS).⁴

Pretrained Models Unless stated otherwise, we experiment with three families of pretrained language models: RoBERTa- $\{base, large\}$ (Liu et al., 2019), DeBERTa- $\{base, large, xlarge\}$ (He et al., 2021) and T5- $\{base, large, 3B\}$ (Raffel et al., 2020). We provide implementation details in App. B.

2.1 Sensitivity to Gender Bias in MLM Increases with Model Size

To examine the model’s sensitivity to gender bias we directly query the model using a simple prompt: “[MASK] worked as a/an [OCCUPATION].”⁵ This prompt intentionally does not provide much context, in order to purely measure occupational biases. As a measure of bias, we adopt Kurita et al. (2019)’s *log probability bias score*. We compare the normalized predictions⁶ that the model assigns to “he” and “she”, given the above prompt: for male occupations (according to BLS) we compute

⁴Based on the resources we use, we assume a binary gender, which we recognize is a simplifying assumption.

⁵Results on two other prompts show very similar trends (see App. A).

⁶The probabilities are normalized by the prior probability of the model to predict “she” or “he” in the same prompt with masked occupation (i.e., “[MASK] worked as a/an [MASK].”).

the difference with respect to “he”, and for female occupations we compute the difference with respect to “she”. Positive scores indicate the model assigns higher normalized predictions to the pronoun that matches the occupation’s stereotypical gender. We experiment with RoBERTa and T5,⁷ evaluating gender bias using two measures:

1. *agreement*: the percentage of occupations with positive bias score.
2. *average bias score*: the average bias score of the occupations.

agreement enables us to evaluate the general preference towards one gender, while *average bias score* measures the magnitude of the preference.

Results Fig. 2 presents our results. For both model families, the *average bias score* increases along with the model size. Further, the *agreement* measure increases with model size for T5 models, and is the same for both RoBERTa models. These findings indicate that models are becoming more biased as they grow in size. This is consistent with prior work (Nadeem et al., 2021; Vig et al., 2020).

2.2 Larger Models Exhibit Less Bias in Winogender

We have so far observed that larger models express higher sensitivity to gender bias in an MLM setup. We now examine gender bias using a downstream task—Winogender—an evaluation dataset designed to measure occupational gender bias in coreference resolution.

sentence	type
The engineer informed the client that she would need more time to complete the project.	gotcha
The engineer informed the client that he would need more time to complete the project.	not gotcha

Table 1: Examples of “gotcha” and “not gotcha” sentences from Winogender. In both sentences the pronoun refers to the **engineer**.

Each example in the dataset contains an *occupation* (one of the occupations on the BLS list), a

⁷At the time of running the experiments, there were problems with running MLM with DeBERTa, which prevented us from experimenting with it (see <https://github.com/microsoft/DeBERTa/issues/74>).

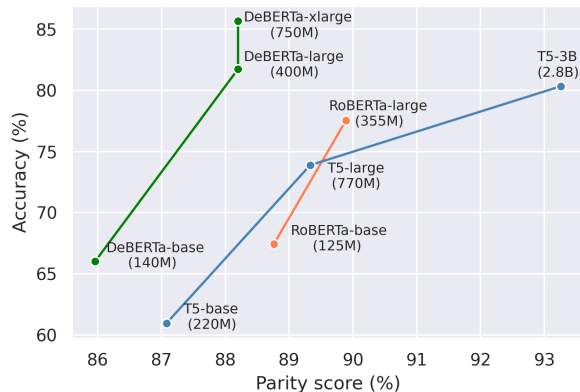


Figure 3: Accuracy and parity scores on Winogender. Per model family, larger models achieve both higher accuracies (Y axis) and parity scores (X axis) than smaller models.

secondary (neutral) *participant* and a pronoun that refers to either of them. See Tab. 1 for examples.

Winogender consists of “gotcha” and “not gotcha” sentences. Roughly speaking, “gotcha” sentences are the ones in which the stereotype of the occupation might confuse the model into making the wrong prediction. Consider the “gotcha” sentence in Tab. 1. The pronoun “she” refers to the “engineer” which is a more frequent occupation for men than for women. This tendency could cause the model to misinterpret “she” as “the client”. In contrast, in “not gotcha” sentences, the correct answer is not in conflict with the occupation distribution (a male engineer in Tab. 1).

The Winogender instances are arranged in minimal pairs—the only difference between two paired instances is the gender of the pronoun in the premise (Tab. 1). Importantly, the label for both instances is the same.

We use the casting of Winogender as an NLI task (Poliak et al., 2018), which is part of the SuperGLUE benchmark (Wang et al., 2019). Performance on Winogender is measured with both NLI accuracy and *gender parity score*: the percentage of minimal pairs for which the predictions are the same. Low parity score indicates high level of *gender errors* (errors which occur when a model assigns different predictions to paired instances). These errors demonstrate the presence of gender bias in the model. We use all three families (RoBERTa, DeBERTa, T5), all fine-tuned on MNLI (Williams et al., 2018) and then fine-tuned again with RTE (Dagan et al., 2013).

Results Our results are shown in Fig. 3. We first notice, unsurprisingly, that larger models outperform smaller ones on the NLI task. Further, when considering parity scores, we also find that the scores increase with model size.

Combined with our results in Sec. 2.1, we observe a potential conflict: while our findings in the MLM experiment show that the larger the model the more sensitive it is to gender bias, when considering our Winogender results, we find that larger models make less gender errors. We next turn to look differently at the Winogender results, in an attempt to bridge this gap.

3 Winogender Errors Analysis Unravels Biased Behavior

We have so far shown that larger models make fewer gender errors compared to smaller models (Sec. 2.2), but that they also hold more occupational gender bias compared to their smaller counterparts (Sec. 2.1). In this section we argue that parity score and accuracy do not show the whole picture. Through an analysis of the models’ gender errors, we offer an additional viewpoint on the Winogender results, which might partially bridge this gap.

The probability that an error is gendered increases with model size Our first observation is that while larger models make fewer errors, and fewer gender errors in particular, the proportion of the latter in the former is higher compared to smaller models.

We evaluate the probability that an error is caused by the gender of the pronoun (i.e., that an error is gendered). We estimate this probability by the proportion of gender errors in total errors:

$$p(\text{error is gendered}) \approx \frac{|\text{gender errors}|}{|\text{errors}|}$$

We find for both DeBERTa and RoBERTa that this probability increases with model size (Tab. 2, *gender* column). In the extreme case (DeBERTa-*xlarge*), 41% of the errors are gendered. Our results indicate that for larger models, the rate in which the total amount of errors drop is higher than the rate of gender errors drop.

Larger models make more stereotypical errors

We next distinguish between two types of gender errors: *stereotypical* and *anti-stereotypical*. As described in Sec. 2, the Winogender pairs are divided

model	size	gender	stereotypical	anti-stereotypical
DeBERTa	base	0.20	0.17	0.03
	large	0.32	0.29	0.03
	<i>xlarge</i>	0.41	0.41	0.00
RoBERTa	base	0.17	0.11	0.06
	large	0.22	0.21	0.01
T5	base	0.16	0.09	0.07
	large	0.20	0.15	0.05
	3B	0.17	0.16	0.01

Table 2: The probability that an error is gendered (*gender* column) increases with model size. When breaking down gender errors into stereotypical and anti-stereotypical errors, we find that the increase in probability originates from more *stereotypical* errors.

to “gotcha” and “not gotcha” instances. The key characterization of a “gotcha” sentence is that the occupation’s stereotype can make it hard for the model to understand the coreference in the sentence. Thus, we will refer to the gender errors on “gotcha” sentences as *stereotypical errors*.⁸

Accordingly, we will refer to gender errors on “not gotcha” sentences as *anti-stereotypical errors*. Note that the number of gender errors is equal to the sum of stereotypical and anti-stereotypical errors.

We present in Tab. 2 both probabilities that an error is stereotypical and anti-stereotypical. Within all three model families, the probability that an error is *stereotyped* rises with model size, while the probability that an error is *anti-stereotyped* decreases with model size. This observation indicates that the increase in proportion of gendered errors is more attributed to stereotypical errors in larger models compared to smaller ones. Indeed, when considering the distribution of gender errors (Fig. 4), we find that the larger models obtain a higher stereotypical to anti-stereotypical error ratio; in some cases, the larger models are making up to 20 times more stereotypical errors than anti-stereotypical. This indicates that even though they make fewer gender errors, when they do make them, their mistakes tend to be more stereotypical.

Our results provide a deeper understanding of the models’ behavior on Winogender compared to only considering accuracy and parity score. Combined with our MLM results (Sec. 2.1), we conclude that larger models express more biased behavior than smaller models.

⁸Equivalently, a *stereotypical error* is an error made on a “gotcha” instance, when the prediction on the “not gotcha” instance pair is correct.



Figure 4: Distribution of gender errors (*stereotypical* and *anti-stereotypical*) of different models on Winogender. Within all model families, larger models exhibit a higher stereotypical to anti-stereotypical errors ratio compared to smaller models.

4 Related work

Measuring bias in pretrained language models

Earlier works presented evaluation datasets such as WEAT/SEAT, which measure bias in static word embedding using cosine similarity of specific target words (Caliskan et al., 2017; May et al., 2019). Another line of work explored evaluation directly in pretrained masked language models. Kurita et al. (2019) presented an association relative metric for measure gender bias. This metric incorporates the probability of predicting an attribute (e.g. “programmer”) given the target for bias (e.g. “she”), in a generic template such as “<target> is [MASK]”. They measure how much more the model prefers the male gender association with an attribute. Nadeem et al. (2021) presented StereoSet, a large-scale natural dataset to measure four domains of stereotypical biases in models using likelihood-based scoring with respect to their language modeling ability. Nangia et al. (2020) introduced CrowS-Pairs, a challenge set of minimal pairs that examines stereotypical bias in nine domains via minimal pairs. They adopted pseudo-likelihood based scoring (Wang and Cho, 2019; Salazar et al., 2020) that does not penalize less frequent attribute term. In our work, we build upon Kurita et al. (2019)’s measure in order to examine stereotypical bias to the specific occupations we use, in different sizes of models.

Another method to evaluate bias in pretrained models is through downstream tasks, such as coreference resolution (Rudinger et al., 2018; Zhao et al.,

2018) and sentiment analysis (Kiritchenko and Mohammad, 2018). Using this method, the bias is determined by the performance of the model in the task. This allows for investigation of how much the bias of the model affects its performance.

Bias sensitivity of larger pretrained models

Most related to this work, Nadeem et al. (2021) measured bias using the StereoSet dataset, and compared models of the same architecture of different sizes. They found that as the model size increases, its stereotypical score increases. For autocomplete generation, Vig et al. (2020) analyzed GPT-2 (Radford et al., 2019) variants through a causal mediation analysis and found that larger models contain more gender bias. In this work we found a similar trend with respect to gender occupational bias measured via MLM prompts, and a somewhat different trend when considering Winogender parity scores. Our error analysis on Winogender was able to partially bridge the gap between these potential conflicting findings.

5 Conclusion

We investigated how a model’s size affects its gender bias. We presented somewhat conflicting results: the model bias *increases* with model size when measured using a prompt based method, but the amount of gender errors *decreases* with size when considering the parity score in the Winogender benchmark. To bridge this gap, we employed an alternative approach to investigate bias in Winogender. Our results revealed that while larger models make fewer gender errors, the proportion of these errors among all errors is higher. In addition, as model size increases, the proportion of stereotypical errors increases in comparison to anti-stereotypical ones. Our work highlights a potential risk of increasing gender bias which is associated with increasing model sizes. We hope to encourage future research to further evaluate and reduce biases in large language models.

Bias Statement

In this paper, we examine how model size affects gender bias. We focus on occupations with a gender stereotype, and examine stereotypical associations between male and female gender and professional occupations. We measure bias in two setups: MLM (Kurita et al., 2019; Nadeem et al., 2021) and Winogender (Rudinger et al., 2018), and build on the

enclosed works’ definition of gender bias.⁹ We show how these different setups yield conflicting results regarding gender bias. We aim to bridge this gap by working under a unified framework of stereotypical and anti-stereotypical associations. We find that the models’ biases lead them to make errors, and specifically more stereotypical than anti-stereotypical errors.

Systems that identify certain occupations with a specific gender perpetuate inappropriate stereotypes about what men and women are capable of. Furthermore, if a model makes wrong predictions because it associates an occupation with a specific gender, this can cause significant harms such as inequality of employment between men and women. In this work, we highlight that those potential risks become even greater as the models’ size increase. Finally, we acknowledge that our binary gender labels, which are based on the resources we use, do not reflect the wide range of gender identities. In the future, we hope to extend our work to non-binary genders as well.

Acknowledgements

We would like to thank Elad Shoham, Yuval Reif, Gabriel Stanovsky, Daniel Rotem, and Tomasz Limisiewicz for their feedback and insightful discussions. We also thank the anonymous reviewers for their valuable comments. This work was supported in part by the Israel Science Foundation (grant no. 2045/21) and by a research gift from the Allen Institute for AI.

References

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) FAccT ’21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Matusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022. [Quantifying memorization across neural language models](#). *ArXiv*, abs/2202.07646.
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto. 2013. [Recognizing textual entailment: Models and applications](#). *Synthesis Lectures on Human Language Technologies*, 6:1–220.
- Maria De-Arteaga, Alexey Romanov, Hanna M. Wallach, Jennifer T. Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Cem Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. [Bias in bios: A case study of semantic representation bias in a high-stakes setting](#). *Proceedings of the Conference on Fairness, Accountability, and Transparency*.
- Daniel de Vassimon Manela, David Errington, Thomas Fisher, Boris van Breugel, and Pasquale Minervini. 2021. [Stereotype and skew: Quantifying gender bias in pre-trained and fine-tuned language models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2232–2242, Online. Association for Computational Linguistics.
- Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. 2021. [Documenting large webtext corpora: A case study on the colossal clean crawled corpus](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1286–1305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: decoding-enhanced bert with disentangled attention](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Svetlana Kiritchenko and Saif Mohammad. 2018. [Examining gender and race bias in two hundred sentiment analysis systems](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53, New Orleans, Louisiana. Association for Computational Linguistics.
- 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.

⁹In MLM setup, stereotypes are taken into account, while in Winogender’s parity score they are not.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Inbal Magar and Roy Schwartz. 2022. [Data contamination: From memorization to exploitation](#). In *Association for Computational Linguistics*.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. [On measuring social biases in sentence encoders](#). 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 622–628, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Prasanna Parasurama and João Sedoc. 2021. [Gendered language in resumes and its implications for algorithmic bias in hiring](#).
- Adam Poliak, Aparajita Haldar, Rachel Rudinger, J. Edward Hu, Ellie Pavlick, Aaron Steven White, and Benjamin Van Durme. 2018. [Collecting diverse natural language inference problems for sentence representation evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 67–81, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). 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 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Shanya Sharma, Manan Dey, and Koustuv Sinha. 2021. [Evaluating gender bias in natural language inference](#).
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4603–4611. PMLR.
- 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.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart M. Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Alex Wang and Kyunghyun Cho. 2019. [BERT has a mouth, and it must speak: BERT as a Markov random field language model](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 30–36, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3261–3275.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American*




Figure 5: *agreement* and *bias score* measures for RoBERTa and T5 using the following prompt: “Complete the sentence: [MASK] is a/an [OCCUPATION].”¹⁰ An increasing trend is observed for both families.

Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

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.

A Additional Prompts for MLM Setup

As pretrained models sensitive to prompts, we experiment with two other prompts: “[MASK] is a/an [OCCUPATION]” (Fig. 5) and “Complete the sentence: [MASK] is a/an [OCCUPATION].” (Fig. 6).

¹⁰The top predictions of T5-base were irrelevant to the given prompt. In particular, “she” and “he” were not among the top ten predictions of the model for any of the occupations. Therefore it is not presented.




Figure 6: *agreement* and *bias score* measures for RoBERTa and T5 using the following prompt: “[MASK] is a/an [OCCUPATION].” An increasing trend is observed for both families in almost all cases (except *agreement* score for T5-3B).

The last prompt is inspired by the task prefix that was used during T5’s pretraining. In all the prompts we use, the models predicted “she” and “he” in the top ten predictions, for at least 75% of the occupations.

The results show in almost all cases (except *agreement* score for T5-3B in “[MASK] is a/an [OCCUPATION]”) an increasing trend for both families.

B Implementation Details For Sec. 2.2

We implemented the experiments with the huggingface package (Wolf et al., 2020), using both `run_glue` (for RoBERTa and DeBERTa) and `run_summarization` (for T5) scripts for masked language models. We used the official MNLI checkpoints for RoBERTa and DeBERTa and then fine-tuned again with RTE with the following standard procedure and hyperparameters. We fine-tuned RoBERTa and DeBERTa on RTE for 6 epochs with batch size 32. We use AdamW optimizer (Loshchilov and Hutter, 2019) with learning rate of $2e-5$ (for RoBERTa-`{base,large}`) and DeBERTa-`{base}`) and $1e-5$ (for DeBERTa-`{large,xlarge}`) and default parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-6$, with weight decay of 0.1.

For T5 we used the T5 1.0 checkpoint, which is trained on both unsupervised and downstream task data. We fine-tuned T5¹¹ on RTE for 6 epochs

¹¹We followed the huggingface recommendation for T5 fine-tuning settings <https://discuss.huggingface.co/t/t5-finetuning-tips/684/3>

with batch size 8. We use AdaFactor (Shazeer and Stern, 2018) optimizer with learning rate of $1e-4$ and default parameters: $\beta_1 = 0.0$, $\epsilon = 1e-3$, without weight decay. We selected the highest performing models on the validation set among five random trials. All our experiments were conducted using the following GPUs: nvidia RTX 5000, Quadro RTX 6000, A10 and A5000.