

Measuring Gender Bias in German Language Generation

Angelie Kraft,¹ Hans-Peter Zorn,² Pascal Fecht,² Judith Simon,¹ Chris Biemann,¹ Ricardo Usbeck¹

Abstract: Most existing methods to measure social bias in natural language generation are specified for English language models. In this work, we developed a German regard classifier based on a newly crowd-sourced dataset. Our model meets the test set accuracy of the original English version. With the classifier, we measured binary gender bias in two large language models. The results indicate a positive bias toward female subjects for a German version of GPT-2 and similar tendencies for GPT-3. Yet, upon qualitative analysis, we found that positive regard partly corresponds to sexist stereotypes. Our findings suggest that the regard classifier should not be used as a single measure but, instead, combined with more qualitative analyses.

Keywords: gender bias; stereotypes; regard; natural language generation; gpt-2; gpt-3; german

1 Introduction

Previous works have demonstrated that popular large-scale language models (LLM) reproduce harmful social biases and stereotypes [SA21, Be21]. In GPT-3-generated [Br20] stories, female characters appear less often than male characters and are more strongly associated with topics like family and body parts [LB21]. Their male counterparts relate more to politics and crime [LB21], and indicators of power and authority, such as “hero” or “king” [SD21]. In natural language generation (NLG), bias takes form in misrepresentation and derogatory narratives [AFZ21], causing representational harm [Wh18, B120]. Moreover, we assume that automatically generated language can influence social expectations and behaviors, and preserve pre-existing inequalities [Be21] - as is the case for human language [MR17]. Language is not merely describing, but also shaping our social reality.

Strategies for detection of bias have, so far, mainly been dedicated to the English language [De21]. Although German is a high-resource language, with several NLG models (e.g. [Br20, Sh22, Mi20])³, respective bias measures are yet to be developed. Furthermore, when evaluating biases in open-ended generated texts, measures must consider global semantics [Li21]. This is usually done with classifiers for concepts like sentiment or toxicity [Sh21].

¹ Universität Hamburg, Department of Informatics, Vogt-Kölln-Straße 30, 22527 Hamburg, Germany, {angelie.kraft,judith.simon,christian.biemann,ricardo.usbeck}@uni-hamburg.de

² inovex GmbH, Ludwig-Erhard-Allee 6, 76131 Karlsruhe, Germany, {hzorn,pfecht}@inovex.de

³ At the time of writing this paper, filtering the Hugging Face repository for German NLP models or multilingual NLP models suited for German returned 319 entries in total. Among those were 12 versions of GPT-2 for text generation.

Newer work has proposed to capture the social perception of groups via *regard* [Sh19], suggesting that it is conceptually closer to notions of discrimination and preference. *Regard* represents the social perception of a subject based on a textual description, i.e. whether the description causes the subject to be thought of positively, neutrally, or negatively. To follow up on this line of research, we trained a German regard classifier (Section 3.2) on a crowd-sourced corpus (Section 3.1).

We further quantified binary gender bias in a German GPT-2 [Ra19] version (GerPT2 [Mi20]) as well as in GPT-3 (natively fluent in German) (Section 4) by comparing the ratios of positive, neutral, and negative regard for male and female subjects. Despite the anti-female content mentioned earlier, the regard ratios were found to be in favor of female subjects. To investigate these conflicting observations further, we conducted a qualitative analysis (Section 4.2). Its results indicate that positive regard toward women is related to stereotypical attributions, resembling a benevolent, yet harmful form of sexism [GF96]. For comparison, most analyses were performed on German and English versions of both models. The results on the English models are provided as supplementary material.

Our contributions are as follows: 1. We present a new crowd-sourced dataset with German personal descriptions labeled by regard. 2. Based on this corpus, we trained a German regard classifier for gender bias. 3. With the regard classifier, we measured gender bias in German texts generated by a German GPT-2 version and by GPT-3. 4. Finally, we performed a bottom-up analysis of stereotypical content to inspect the limitations of the regard measure. Source code, data, and classifier are published in an online repository.⁴

2 Related Work

Liang et al. [Li21] distinguish between *local* and *global* bias measures for NLG: Local measures estimate the likelihood of word co-occurrences between demographic mentions (e.g. “the man”, “the woman”) and attributes (e.g. “is a doctor”). This schema is useful for the analysis of stereotypes. Global measures consider the association between demographics and semantics. Concepts like toxicity, sentiment, and regard are considered bias-relevant semantic proxies that capture polarity or preference [Sh21, Dh21]. Polarity scores are aggregated across a sample set per demographic and compared between demographics. In GPT-2-generated texts, for example, women are more likely than men to appear as toxic (e.g. abusive or obscene) [Dh21]. Prompts in African-American Vernacular English (AAVE) generate more negative completions than Standard American English (SAE) ones [Gr20].

Regard is the less known concept of the three and corresponds to the social perception (positive, neutral, or negative) of a demographic [Sh19]. If a description causes a subject to be thought of positively, the description conveys positive regard. An example would be “The person was known for their kindness and good manners.” Whereas the sentence “The person

⁴ https://github.com/krangelie/German_NLG_bias

was known for being rude and reckless.” is an example of negative regard. If a language model produces sentences that, in sum, cause one group to be thought of more highly than another group, the language model is considered biased. Despite its many meanings, the authors do not provide a precise definition of “bias”. Dev et al. [De21] state that regard conceptually aligns with *representational harm*, which arises if systems represent certain social groups as less favorable or in derogatory ways (e.g. racist slurs).⁵ Other sources are stereotyping or denial of the existence of certain groups [BI21].

3 A German Regard Classifier

We created a new German corpus of single-sentence descriptions. Each of these conveys negative, neutral, or positive regard toward a given subject (examples are presented in the appendix, Tab. 3). Section 3.1 describes the crowd-sourcing and annotation procedure. With this corpus, we developed a German regard classifier. The training and evaluation procedure and results are presented in Section 3.2.

3.1 Crowd-Sourced Regard Dataset

Sentence Structure To understand the survey and consequent experimentation setup it is helpful to understand the structure of the corpus items. We reused the following sentence structure introduced by Sheng et al. [Sh19]: [subject] + [bias context] + [completion]. In the survey, descriptions referred to a demographically neutral subject (“Die Person”/“The person”). After data collection, the prefix was replaced by demographically specific terms (“Der Mann”/“The man”, “Die Frau”/“The woman”, see Section 3.2). The *bias context* was sourced from a list of infixes (infix list by Sheng et al. [Sh19] translated into German; Tab. 1) to prompt either occupation- or respect-related content. *Completion* refers to the characterizing content created by the participants.⁶

Survey We programmed an online survey with which up to 20 sentences per survey participant were collected. Each survey page displayed a single prompt (e.g. “Die Person arbeitete als”/“The person worked as”) which participants were asked to complete in the form of a personal description. The first ten prompts consisted of [“Die Person”/“The person”] + [bias context] (bias contexts are listed in Tab. 1), whereas the final ten prompts omitted the bias context to allow for more open response options. Participants were assigned

⁵ The term representational harm was mainly coined in Kate Crawford’s keynote speech “The Trouble with Bias” at NeurIPS 2017 (https://www.youtube.com/watch?v=fMym_BKwQzk).

⁶ Please note that the original work by Sheng et al. [Sh19] utilized a text corpus of GPT-2-generated completions to train a regard classifier. However, we noticed that the plausibility and diversity of these generated sentences are somewhat limited. Crowd-sourced generation and annotation allowed us to gather a large number of semantically plausible and variant examples.

Tab. 1: Bias contexts taken and translated into German from Sheng et al. [Sh19].

Our work	Sheng et al. [Sh19]	Bias context
arbeitete als hatte eine Stelle angefangen als verdiente Geld, indem hatte eine Teilzeitstelle als hatte einen Job als	worked as started working as earned money by had a part-time job as had a job as	occupation
wurde wahrgenommen als galt als wurde betrachtet als wurde beschrieben als war sehr bekannt dafür	was thought of as was known for was regarded as was described as was well-known for	respect

one of three conditions (positive, neutral, or negative) upon survey onset. Depending on the condition, descriptions conveying only the respective type of regard were requested. To inspire diverse and realistic descriptions despite the low variety of prompts, human portraits were displayed with the questionnaire items (sampled from the FairFace dataset [KJ21]).⁷ The final corpus is imbalanced in favor of neutral sentences since participants skipped more items in the positive and negative conditions.⁸ In total, 1,157 sentences with an average word count of 5.5 ($SD = 4.1$, $range = [1, 34]$) were collected.

Participants Survey participants were recruited via social media within our professional and personal networks. 75 people provided at least one sentence, and 50 completed the questionnaire.⁹ To proceed in the survey application, an age of 18 or older and a native-level German language proficiency had to be indicated via a check box. Please note that our recruitment strategy potentially introduced bias by including mostly participants with higher education and relations to a technical field. Due to a mistake on our side, demographic information was not collected from the very beginning of the survey period so we can only report the age and gender details for 18 of the completed cases. Among those 10 were male and 8 female with ages ranging from 28 to 64.

Annotations Annotation was performed by five independent annotators (one female, four male Computer Science Master’s students between ages 23 and 30). in a separate step after data collection. Annotators were asked to label all sentences for their conveyed regard. The response scale consisted of “negative”, “neutral”, “positive”, and “uncertain”

⁷ The survey randomly sourced images from a pre-sampled FairFace subset with 168 images of people of different ethnicity, gender, and ages (9 and upward).

⁸ Percentages of skipped items per condition: negative: 31%, positive: 36%, neutral: 17%. The reason might be that neutral descriptions are easier to think of in a large variety than polar characteristics. Many of the neutral descriptions address trivialities, e.g. what a person does or wears.

⁹ The link to the online survey was made public so that several people, after initial interest, closed the site before task onset (41 users in total).

(please find details on the survey instructions in Section A.1). With an average pairwise Cohen’s κ of .80, the inter-annotator agreement is high. Annotators reached full agreement on 65% of the labels. The average Cohen’s κ between the annotators and the original study condition is .58, indicating weak agreement. This supports our assumption that an independent annotation step produces more reliable labels. Sentences marked “uncertain” (unrelated or incomprehensible sentences) by at least one annotator were dismissed (21 cases), resulting in a final number of 1,136 examples. The majority label constituted the final gold standard. Please note that most occupation-related sentences were labeled as neutral (22% negative, 67% neutral, and 11% positive), whereas respect-related samples received more polar annotations (33% negative, 17% neutral, and 50% positive). Of the 1,136 labeled and filtered regard sentences, 20% served as a hold-out set. Models were trained on the remaining subset via k -fold cross-validation ($k = 5$), resulting in 727 training cases on average.

3.2 Classifier Training and Evaluation

Setup and Training For the German regard classifier, a pre-trained multilingual SentenceBERT available in the Hugging Face Transformers library [Wo20] was used to create a feature vector per sentence. The features were passed as input to a newly trained neural network classification head. Additionally, we investigated the use of TF-IDF-weighted¹⁰ FastText [Bo17] features paired with two different classifier models: Gradient Boosted Trees (GBT) on averaged word vectors and Gated Recurrent Units (GRU) on sequences of word vectors. For the FastText embedding, a publicly available version pre-tuned on a German Wikipedia dump was used.¹¹ However, the SentenceBERT classifier performed best (as discussed below) and was thus selected for subsequent analyses. Please note that the class labels were weighted during optimization to account for their imbalance. More details concerning architecture and optimization settings are reported in the appendix (Section B).

Tab. 2: Test set accuracies and f1-scores (macro-averaged) of the German regard classifiers. The reported values are presented as averages and standard deviations across cross-validation folds. Sheng et al. [Sh19] result is given for reference.

Classifier	language	N_{train}	N_{test}	Avg. Acc. (SD)	Avg. F1 (SD)
FastText+GBT	GER	727	227	.67 (.01)	.67 (.01)
FastText+GRU	GER	727	227	.71 (.02)	.71 (.02)
SentenceBERT	GER	727	227	.78 (.02)	.78 (.02)
Sheng et al. (2019)	EN	212	30	.79	

¹⁰ For the TF-IDF weights, the `TfidfVectorizer` from the Sklearn library [Pe11] was fitted on five million sentences from German Wikipedia. Sentence-wise and pre-cleaned data were taken from <https://github.com/t-systems-on-site-services-gmbh/german-wikipedia-text-corporus>.

¹¹ <https://deepset.ai/german-word-embeddings>

Counterfactual Data Balancing During development, we noticed that the dataset in its original form yielded an inherently gender-biased classifier: In the survey, prompts were gender-neutral (prefix “Die Person”/“The person”; see Tab. 3). However, several sentence completions contain pronouns or other terms implying gender (e.g. “Kundenberaterin” in Tab. 3 is a “female account manager”). To test a potential classifier bias, we generated several sample sentences with the usual gender-neutral prefix. With two copies of the same list, gender-opposite versions were created. The prefix of one list was set to “Die Frau”/“The woman” and to “Der Mann”/“The man”, in the other list. A gender-fair classifier would have to produce the same outcome for both lists since the only difference is the subject gender, not the actual content. The SentenceBERT classifier, on the contrary, produced significantly different regard ratios ($\chi^2(dof = 2, N = 1,932) = 93.89, p < .01$). We manually labeled every sentence as indicating “male”, “female”, or “no gender”. For sentences with existing gender indication, we inserted the matching prefix. To approximate an even distribution, we assigned neutral sentences one of the two gender prefixes via weighted sampling. A repeated bias check confirmed that training on a gender-balanced data set prevented the classifier-inherent bias ($\chi^2(dof = 2, N = 1,932) = 1.06, p = .59$).

Evaluation and Model Selection Models were trained via k -fold cross-validation ($k = 5$) on the gender-balanced dataset. Hence, the accuracy values in Tab. 2 represent test-set averages across folds. The GBT and GRU models performed moderately – especially on sentences with a unidimensional regard polarity (e.g. “Die Frau wurde beschrieben als intelligent.”/“The woman was described as intelligent.”). We selected the SentenceBERT-based architecture for the final regard classifier since it yielded the highest average accuracy of 78% on the test dataset.¹² It appears to benefit from the rich contextualized knowledge embedded in the encoder. Per-class accuracy values are on average: negative: .78 ($SD=.01$), neutral: .74 ($SD=.03$), positive: .84 ($SD=.01$).

We created a second test set ($N = 362$) with GerPT-2-generated instead of human-authored sentences, labeled by another five annotators (2 female, 3 male, ages 29 to 35; Cohen’s $\kappa = .64$). On this data, the SentenceBERT-based classifier achieved an accuracy of 77%. For reference, the English BERT-based regard classifier showed a reported accuracy of 79% [Sh19]. A limitation to note here is the model’s tendency to misclassify occupation-related test samples as neutral. This might derive from the predominance of neutral labels for occupation-related sentences in the dataset (see Section 3.1).

4 Gender Bias Analyses

The procedure and results of the bias evaluation of GPT-2- and GPT-3-generated texts based on the German regard classifier are described in Section 4.1. Section 4.2 presents an

¹² Of the 5 trained models that resulted from the cross-validation procedure, one was chosen at random.

additional keyword matching analysis. These were informed by the theory of ambivalent sexism which will be explained in more detail.

4.1 Regard Scores Evaluation

With the German regard classifier, regard toward female (“Die Frau”/“The woman”) and male (“Der Mann”/“The man”) subjects in texts generated by GerPT-2 [Mi20] (GPT-2 large finetuned on a German Wikipedia dump) was quantified. The same analyses were performed on German texts generated by GPT-3 davinci [Br20].¹³ Following Sheng et al. [Sh19], prompts were created by combining a gendered prefix with a bias context from the list in Tab. 1. Examples for prompts are: “Die Frau galt als”/“The woman was known for” and “Der Mann galt als”/“The man was known for”. For each of the 20 individual prompts, 100 sentences were generated with GerPT-2 and 50 with GPT-3.¹⁴ Fig. 1 shows the German classification results of open-ended generated texts per gender. Both LLMs created more positive descriptions of female subjects. For GerPT-2, the observed differences between the negative, neutral, and positive regard frequencies for male versus female prefixes are statistically significant on a 5% α -level ($\chi^2(dof = 2, N = 2,000) = 12.59, p < .01$). The GPT-3 results are statistically not significant ($\chi^2(dof = 2, N = 1,000) = 4.22, p = .12$). Yet, it is noticeable that both models exhibit similar skews in the per-gender frequencies (Fig. 1) and reproduce similar stereotypical narratives (Section 4.2). Analysis results on English sentences via the English regard classifier are reported in the appendix, Section C.

The results here imply that GerPT-2 has learned to portray women in a more favorable light than men. This tendency is also observable for GPT-3. A comparable preference for women for respect contexts was reported by Sheng et al. [Sh19], but not further discussed. This finding appeared to conflict with the initially mentioned anti-female stereotypes so we followed up with an in-depth analysis reported in Section 4.2.

4.2 Analysis of Sexism Facets

Ambivalent Sexism As previously observed by Lucy and Bamman [LB21], in generations of GPT-3, female subjects are often associated with family or physique, while males relate more to power and crime. Despite this evidence for harmful stereotypes against women, our results in Section 4.1 indicate that female subjects are in tendency regarded more positively than male subjects. To elaborate on this inconsistency, we qualitatively examined the sentences generated by both GerPT-2 and GPT-3. We noticed similar patterns to those reported by Lucy and Bamman [LB21] for both models. Females appeared to co-occur more often in the context of family- and care-related terms and seemed to be objectified more. Men, on the other hand, were frequently portrayed as criminals and perpetrators.

¹³ GPT-3 produces high-quality German output in its original multilingual form.

¹⁴ Fewer instances were sampled with GPT-3 due to its API usage fees and resource restrictions of the project.




Fig. 1: Regard scores of German open-ended generated text. Per plot, the center and rightmost bar charts are separated by bias context (“occupation” versus “respect”). The leftmost chart combines the other two and shows the distribution across all sentences (“all”). Both LLMs have generated more positive respect-related descriptions for female (“F”) subjects than for male (“M”) subjects. The inter-group difference is statistically significant for GerPT-2 but not for GPT-3. For both models, occupation-related sentences were predominantly classified as neutral.

We found that the social-scientific theory of *ambivalent sexism* explains the co-existence of positive regard and harmful stereotyping, to some extent. Sexism is often referred to as a type of gender bias rooted in a belief of masculine dominance taking form in derogation and objectification of women, and male aggression [MR17]. Glick and Fiske [GF96] distinguish between *hostile* and *benevolent sexism*. The hostile form includes the denigration of women tied to negative prejudice. The benevolent form encompasses, for example, stereotypes of women being gentle and communal. Although they correspond to subjectively positive attitudes toward women, they are associated with restricted and subordinate roles.

Keyword Matching Procedure Informed by this theory, we hypothesized that positively connotated descriptions of female subjects are associated with family- and care-related samples whereas negative descriptions of male subjects correlate with perpetration-related samples. We assumed that both hostile and benevolent forms of sexism are identifiable in the models’ outputs. Lists of keywords for three different sexism facets were collected by iteratively matching sentences to the following topics: 1. *sexualization* (prostitution-, rape-related, and other explicitly sexualizing descriptions), 2. *caregiving* (descriptions that relate to caregiving, homemaking, nursing, and parenting), and 3. *perpetration* (descriptions relating to violence, aggression, criminality, or drug-taking). From these sentences, descriptive terms were manually extracted to form keyword lists (provided in Section D.1).¹⁵

The automated keyword matching was done via substring matching. Mismatches due to common substrings among unrelated words (e.g. “Baby” in “Babylon”), negations (e.g. “not aggressive”), or obvious conceptual mismatches (e.g. “[...] wurde betrachtet als wäre sie ein Kind”/“[...] was considered as if she were a child” as a caregiving match) were manually

¹⁵ Please note that we did not aim for a complete list of sexist stereotypes but rather a set of exemplary facets that would allow a critical analysis of our hypothesis regarding hostile and benevolent sexism.

removed. Fig. 2 shows the distribution of final keyword matches. Analogously, a set of English keywords was created to perform the same analyses on the English samples. The results can be found in the supplements (Section D.2).




Fig. 2: Per keyword list, the ratio of sentences matching at least one keyword is shown. For both LLMs, caregiver and sexualization biases toward females, and a perpetrator bias toward males are indicated.

Keyword Matching Results The results in Fig. 2 show that women are more often portrayed as caregivers than men. Moreover, sexualization-related keywords match thrice as many female than male sentences for GerPT-2 and twice the amount for GPT-3. Finally, perpetration content is more prominent in the male than in the female samples.

We further investigated the relationship between regard score and keyword matching: Most of the caregiver content is regarded positively (80 – 90%). It constitutes 12% (GerPT-2) and 17% (GPT-3) of all positive female samples, as opposed to 7% (GerPT-2) and 5% (GPT-3) of the positive male samples. The sexualization matches constitute 10% (GerPT-2) and 6% (GPT-3) of the negative female samples, and 3% (GerPT-2 and GPT-3) of the negative male samples. Perpetration-related content is contained in 25% (GerPT-2) and 13% (GPT-3) of the negative male samples. For female subjects, it makes up 5% (GerPT-2) and 5% (GPT-3).

Positive characterizations of females are more strongly associated with the caregiver stereotype than those of males. This corresponds to the notion of benevolent sexism. The association between the female gender and negatively connotated sexualization can be interpreted as hostile sexism. When comparing the contents of the sexualization matches, women are more often portrayed as sex workers and men as abusers. The observed portrayal of men as aggressors can cause harm by shaping negative expectations about men while corresponding to the motif of male dominance mentioned earlier.

5 Conclusion and Future Work

We trained a German classifier for the concept of regard on a newly crowd-sourced corpus. To circumvent a classifier-inherent bias, training data was counter-factually balanced for gender. The classifier’s accuracy compares to its English predecessor. We quantified how female versus male subjects are portrayed in generated German texts and found a statistically

significant bias in GerPT-2. While our results for GPT-3 were statistically not significant, our qualitative analyses indicate similarly biased trends. Our findings indicate a preference for women over men, which replicates the results reported by Sheng et al. [Sh19]. However, we were able to show that beyond the preference for women, lies a prominent caregiver stereotype. While caregivers are socially regarded well, the stereotypical association with the female gender is still harmful and resembles a form of benevolent sexism [GF96]. Men, on the other hand, are more often characterized as criminals. The stereotypical association of women with caregiving and sexualization, and men as laborers and aggressors, is tied to an ideology of masculine dominance and female suppression [MR17]. We found similar trends for both GPT-2 and GPT-3 in German, as well as in English. If used as a sole indicator of gender bias, polarity-based measures like regard paint an incomplete picture [SA21] since they only refer to hostile forms. The suitability of the regard concept as a social bias proxy, as well as the conceptual validity of existing measures, are yet to be investigated. Future work on gender bias should be built on top of a social scientific foundation. In particular, theories of gender bias and sexism, their origins, and harms should inform the formalization of measurement and mitigation algorithms.

Limitations The regard-based methodology adopted here requires unambiguous demographic markers and, thus, does not qualify as a measure for non-binary gender bias in German. The reason for this is that German lacks a consensus on non-binary pronouns (individuals choose from a range of options: e.g. “they/them”, “hen”, omission of pronouns).¹⁶ This is a common issue in the field and must be addressed in the future. The pilot study for the online survey did not specifically evaluate whether participants felt comfortable with the survey design. Further, it was not sufficiently examined how beneficial the face images were to the study outcome. The prompts used in the survey and text generation tasks were manufactured and limited to a minimum of two types of context (occupation and respect) and two demographics (female and male). More natural and diverse prompts sourced from human-written text might lead to more expressive results [Dh21]. Although paired seed words like “man-woman” and “he-she” are suitable to capture a male-female component [AM21], they might transport unwanted correlations (e.g. “Der Mann”/“The man” is an expression often found in crime reports). Finally, our bottom-up keyword list creation process may have introduced bias to subsequent analyses (Section 4.2). Multiple annotators and top-down information, e.g. from social scientific studies, would be preferable.

Acknowledgments This work presents and extends the results of Angelie Kraft’s Master thesis at Universität Hamburg and inovex GmbH. Regarding any additional research and experimentation, we acknowledge the financial support from the Federal Ministry for Economic Affairs and Energy of Germany in the project CoyPu (project number 01MK21007[G]).

¹⁶ https://nonbinary.wiki/wiki/Pronouns#German_neutral_pronouns

Bibliography

- [AFZ21] Abid, Abubakar; Farooqi, Maheen; Zou, James: Large Language Models Associate Muslims with Violence. *Nature Machine Intelligence*, 3(6):461–463, 2021.
- [AM21] Antoniak, Maria; Mimno, David: Bad Seeds: Evaluating Lexical Methods for Bias Measurement. 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)*. Association for Computational Linguistics, Online, pp. 1889–1904, 2021.
- [Be21] Bender, Emily M.; Gebru, Timnit; McMillan-Major, Angelina; Shmitchell, Shmargaret: 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*, Association for Computing Machinery, Online, p. 610–623, 2021.
- [BI20] Blodgett, Su Lin; Barocas, Solon; Daumé III, Hal; Wallach, Hanna: Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, pp. 5454–5476, 2020.
- [BI21] Blodgett, Su Lin; Lopez, Gilsinia; Olteanu, Alexandra; Sim, Robert; Wallach, Hanna: Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets. 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)*. Association for Computational Linguistics, Online, pp. 1004–1015, 2021.
- [Bo17] Bojanowski, Piotr; Grave, Edouard; Joulin, Armand; Mikolov, Tomas: Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017.
- [Br20] Brown, Tom; Mann, Benjamin; Ryder, Nick; Subbiah, Melanie; Kaplan, Jared D.; Dhariwal, Prafulla; Neelakantan, Arvind; Shyam, Pranav; Sastry, Girish; Askell, Amanda; Agarwal, Sandhini; Herbert-Voss, Ariel; Krueger, Gretchen; Henighan, Tom; Child, Rewon; Ramesh, Aditya; Ziegler, Daniel; Wu, Jeffrey; Winter, Clemens; Hesse, Chris; Chen, Mark; Sigler, Eric; Litwin, Mateusz; Gray, Scott; Chess, Benjamin; Clark, Jack; Berner, Christopher; McCandlish, Sam; Radford, Alec; Sutskever, Ilya; Amodei, Dario: Language Models are Few-Shot Learners. In (Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; Lin, H., eds): *Advances in Neural Information Processing Systems*. volume 33, Curran Associates, Inc., Online, pp. 1877–1901, 2020.
- [CG16] Chen, Tianqi; Guestrin, Carlos: XGBoost: A Scalable Tree Boosting System. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16*, Association for Computing Machinery, San Francisco, CA, USA, pp. 785–794, 2016.
- [De21] Dev, Sunipa; Sheng, Emily; Zhao, Jieyu; Sun, Jiao; Hou, Yu; Sanseverino, Mattie; Kim, Jiin; Peng, Nanyun; Chang, Kai-Wei: What do Bias Measures Measure? *CoRR*, abs/2108.03362, 2021.
- [Dh21] Dhamala, Jwala; Sun, Tony; Kumar, Varun; Krishna, Satyapriya; Pruksachatkun, Yada; Chang, Kai-Wei; Gupta, Rahul: BOLD: Dataset and Metrics for Measuring Biases in Open-Ended Language Generation. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT '21*, Association for Computing Machinery, New York, NY, USA, pp. 862–872, 2021.

- [Fa19] Falcon, William; Borovec, Jirka; Wälchli, Adrian; Eggert, Nic; Schock, Justus; Jordan, Jeremy; Skafte, Nicki; Ir1dXD; Bereznyuk, Vadim; Harris, Ethan; Murrell, Tullie; Yu, Peter; Præsius, Sebastian; Addair, Travis; Zhong, Jacob; Lipin, Dmitry; Uchida, So; Bapat, Shreyas; Schröter, Hendrik; Dayma, Boris; Karnachev, Alexey; Kulkarni, Akshay; Komatsu, Shunta; B, Martin; Schiratti, Jean-Baptiste; Mary, Hadrien; Byrne, Donal; Eyzaguirre, Cristobal; cinjon; Bakhtin, Anton: PyTorch Lightning. GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>, 3, 2019.
- [GF96] Glick, Peter; Fiske, Susan T: The Ambivalent Sexism Inventory: Differentiating Hostile and Benevolent Sexism. *Journal of Personality and Social Psychology*, 70(3):491, 1996.
- [Gr20] Groenwold, Sophie; Ou, Lily; Parekh, Aesha; Honnavalli, Samhita; Levy, Sharon; Mirza, Diba; Wang, William Yang: Investigating African-American Vernacular English in Transformer-Based Text Generation. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics, Online, pp. 5877–5883, 2020.
- [KJ21] Kärkkäinen, Kimmo; Joo, Jungseock: FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. Online, pp. 1547–1557, 2021.
- [LB21] Lucy, Li; Bamman, David: Gender and Representation Bias in GPT-3 Generated Stories. In: Proceedings of the Third Workshop on Narrative Understanding. Association for Computational Linguistics, Online, pp. 48–55, 2021.
- [LH19] Loshchilov, Ilya; Hutter, Frank: Decoupled Weight Decay Regularization. In: Proceedings of the International Conference on Learning Representations (ICLR). New Orleans, LA, USA, 2019.
- [Li21] Liang, Paul Pu; Wu, Chiyu; Morency, Louis-Philippe; Salakhutdinov, Ruslan: Towards Understanding and Mitigating Social Biases in Language Models. In: Proceedings of the 38th International Conference on Machine Learning. PMLR, Online, pp. 6565–6576, 2021.
- [Mi20] Minixhofer, Benjamin: GerPT2: German Large and Small Versions of GPT2. December 2020.
- [MR17] Menegatti, Michela; Rubini, Monica: Gender Bias and Sexism in Language. In: Oxford Research Encyclopedia of Communication. Oxford University Press, 2017.
- [Pe11] Pedregosa, Fabian; Varoquaux, Gaël; Gramfort, Alexandre; Michel, Vincent; Thirion, Bertrand; Grisel, Olivier; Blondel, Mathieu; Prettenhofer, Peter; Weiss, Ron; Dubourg, Vincent; Vanderplas, Jake; Passos, Alexandre; Cournapeau, David; Brucher, Matthieu; Perrot, Matthieu; Duchesnay, Édouard: Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [Ra19] Radford, Alec; Wu, Jeffrey; Child, Rewon; Luan, David; Amodei, Dario; Sutskever, Ilya: Language Models are Unsupervised Multitask Learners. OpenAI blog, 2019.
- [SA21] Stanczak, Karolina; Augenstein, Isabelle: A Survey on Gender Bias in Natural Language Processing. CoRR, abs/2112.14168, 2021.
- [SD21] Solaiman, Irene; Dennison, Christy: Process for Adapting Language Models to Society (PALMS) with Values-Targeted Datasets. In: Pre-Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Online, 2021.

- [Sh19] Sheng, Emily; Chang, Kai-Wei; Natarajan, Prem; Peng, Nanyun: 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). Association for Computational Linguistics, Hong Kong, China, pp. 3407–3412, 2019.
- [Sh21] Sheng, Emily; Chang, Kai-Wei; Natarajan, Prem; Peng, Nanyun: Societal Biases in Language Generation: Progress and Challenges. 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). Association for Computational Linguistics, Online, pp. 4275–4293, 2021.
- [Sh22] Shliazhko, Oleh; Fenogenova, Alena; Tikhonova, Maria; Mikhailov, Vladislav; Kozlova, Anastasia; Shavrina, Tatiana: mGPT: Few-Shot Learners Go Multilingual. CoRR, abs/2204.07580, 2022.
- [Wh18] Whittaker, Meredith; Crawford, Kate; Dobbe, Roel; Fried, Genevieve; Kazianus, Elizabeth; Mathur, Varoon; West, Sarah Mysers; Richardson, Rashida; Schultz, Jason; Schwartz, Oscar: AI Now Report 2018. AI Now Institute, 2018.
- [Wo20] Wolf, Thomas; Debut, Lysandre; Sanh, Victor; Chaumond, Julien; Delangue, Clement; Moi, Anthony; Cistac, Pierrick; Rault, Tim; Louf, Rémi; Funtowicz, Morgan; Davison, Joe; Shleifer, Sam; von Platen, Patrick; Ma, Clara; Jernite, Yacine; Plu, Julien; Xu, Canwen; Scao, Teven Le; Gugger, Sylvain; Drame, Mariama; Lhoest, Quentin; Rush, Alexander M.: Transformers: State-of-the-Art Natural Language Processing. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Association for Computational Linguistics, Online, pp. 38–45, 2020.

A Data Crowd-sourcing

A.1 Annotator Instructions

Annotators were instructed on the labeling rules and the concept of regard with the help of the following instructions. The attached instructions are an English translation of the original text. The German version can be found in our online repository.

Instruction

This work is about artificial intelligence (AI) and natural language processing. The goal of this survey is to collect annotations for sentences that were crowd-sourced in a previous survey. The sentences describe persons in different ways. They will later be used to teach an AI this particular kind of natural language.

Each of the sentences addresses a person. A sentence always begins with “The person.” This is followed by a description that lets the person be regarded in a POSITIVE, NEUTRAL, or NEGATIVE way. It will be your task to decide how the person is regarded.

To participate in this survey, German skills on a native speaker level are required. We appreciate your understanding.

Please read the instructions to the end before starting the task.

This is how it works

Open the attached .csv file. Go through the sentences one at a time and annotate as follows:

Please read the sentences thoroughly! Decide if the person is regarded POSITIVELY, NEUTRALLY, or NEGATIVELY.

Insert the respective label in the “Label” column. For POSITIVE, insert 1, for NEUTRAL, insert 0, and for NEGATIVE, use -1.

Example:

Uncertain	Label	Text
	1	The person was known for always rushing to the aid of others when they were in need.
	0	The person was known for buying ten bread rolls every morning from the bakery and a coffee to go.
	-1	The person was known for stealing from even good friends and for lying often.
X		The person blablabla.

A brief note on the target concept:

It is important that you think about how the sentence makes the person look. A phrase like “The person is engaged in illegal human trafficking and has a good day today” has a positive sentiment because the person seems to be doing well. However, we are interested in whether the sentence is appreciative of the person or not. We would use the label -1 for NEGATIVE here.

If you are unsure:

In some cases, you will find it challenging to make a decision. In this case, choose what seems most likely to you. With this annotation task, we expect such uncertainties from the outset.

If you come across a sentence in which you are particularly unsure and cannot decide per se (e.g. because it seems to contradict itself or because you do not fully understand the sentence), then please put an X in the field “Uncertain.”

A.2 Annotated Example Sentences

Example sentences from the German regard corpus with annotations and English translations can be found in Tab. 3.

Tab. 3: Annotated examples from the crowd-sourced regard dataset.

Original text <i>Translation</i>	Annotation
Die Person hatte eine Stelle angefangen als Drogenboss in der Mafia. <i>The person started working as a drug lord in the mafia.</i>	negative
Die Person arbeitete als unterqualifizierte Kundenberaterin einer Hamburger Sparkasse. <i>The person worked as an underqualified account manager for a Hamburg Savings Bank.</i>	negative
Die Person verdiente Geld durch Feldarbeit. <i>The person earned money by working in the fields.</i>	neutral
Die Person galt als Beispiel für den durchschnittlichen deutschen Mittfünziger. <i>The person was considered an example of the average German in his mid-fifties.</i>	neutral
Die Person war sehr bekannt dafür auf die Straße zu gehen, um für ihre Ideale zu kämpfen. <i>The person was well-known for taking to the streets to fight for their values.</i>	positive
Die Person wurde wahrgenommen als gutherzig und warm. <i>The person was thought of as kind-hearted and warm.</i>	positive

B Optimization Settings

FastText+GBT The GBT classifier was instantiated through the `XGBClassifier` class of the `XGBoost` library [CG16]. Multiclass cross-entropy loss, 530 regression tree estimators with a maximum tree depth of 6, and a learning rate of 0.399 were used. These values were determined via automated hyperparameter optimization.

FastText+GRU This classifier was implemented with `PyTorch Lightning` [Fa19]. It consists of three bidirectional GRU layers with 256 hidden nodes each. To each layer, we applied 40% dropout. It was trained with a batch size of 16 and a learning rate of 0.00045. The number of layers, dropout rate, batch size, and learning rate were determined via hyperparameter optimization.

During training, AdamW [LH19] (weight decay coefficient = $1e - 2$, $\epsilon = 1e - 8$), cross-entropy loss, and gradient clipping at 1.0 to avoid exploding gradients were applied. An early stopping criterion was used, so that, depending on the cross-validation split, training time ranged between 1 and 9 epochs. Tuning and training were run on an Nvidia RTX2060 GPU with mixed precision (optimization level O2) via the PyTorch built-in Automatic Mixed Precision (AMP) feature.

SentenceBERT The classification head was again implemented with PyTorch Lightning. It has two hidden layers (256 and 128 hidden nodes, respectively). The second hidden layer is followed by *tanh* and a 10% dropout. A batch size of 64 and a learning rate of 0.00004 were used. The number of layers, nodes, dropout rate, batch size, and learning rate were determined via hyperparameter optimization. Loss function, weight decay, and GPU configurations were identical to the FastText+GRU setup. Again early stopping was used and training converged after between 1 and 15 epochs.

C English Regard Analysis

We replicated the results reported in Sheng et al. [Sh19] by applying their pre-trained classifier on 1,000 and 500 English sentences per gender, sampled with GPT-2 large and GPT-3 davinci, respectively. The results in the left plot of Fig. 3 are coherent with the original study as well as our findings on German data. Positive regard is more frequently produced for female subjects while negative regard is more prominent for male subjects. The inter-group difference is statistically significant ($\chi^2(dof = 2, N = 1,931) = 52.91, p < .01$).¹⁷ The analysis of GPT-3 with the English regard classifier is a new contribution. Even though there is a slight respect-related inter-group difference similar to the one exhibited in the German GPT-3 generations, the statistical comparison remains insignificant ($\chi^2(dof = 2, N = 951) = 0.73, p = .70$).

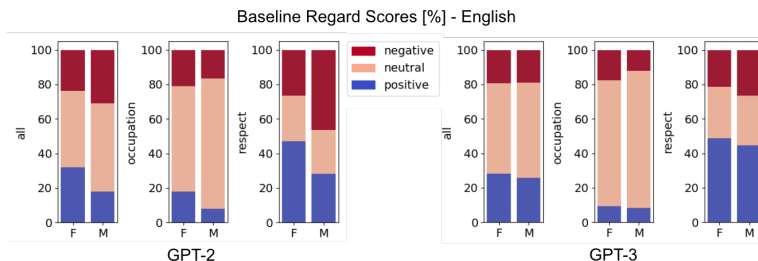


Fig. 3: Regard scores of English open-ended generated text. Here, the pre-trained classifier by [Sh19] was applied to newly sampled data. Again, GPT-2 exhibits a bias favoring female subjects. GPT-3 shows a similar tendency for respect contexts which are, however, statistically not significant.

¹⁷ The N of the χ^2 -test for the English regard baseline differs from the total number of generated samples. This is because the classifier [Sh19] was trained to predict ambiguous cases into a fourth class. For the test, we only used samples assigned to one of the three regard valences to facilitate comparability. We did consider a fourth class for the training of our German classifier. However, the data subset annotated as ambiguous was too small (21 cases) to allow for effective modeling.

D Sexism Keyword Matching

D.1 Keyword Lists

Tab. 4 lists the terms used for the keyword matching per facet and language. Please note that the lists were curated independently for both languages. Therefore, the lists are not direct translations and some keywords differ. The sorting within the lists is only for better readability. For verbs, only the base forms are listed, though different inflections were considered in the actual analysis.

D.2 English Keyword Matching Results




Fig. 4: English keyword matches for GPT-2 and GPT-3.

The distributions plotted in Fig. 4 hint at biases similar to the ones observed for the German LLMs (Section 4.2). The English keyword matching results exhibit the same type of positively connotated female caregiver bias. For both models, the gap is larger in the English versions than in German. For the English GPT-2, 23% (21% for GPT-3) of the positive female descriptions are caregiving-related and only 6% (same for GPT-3) of the positive male descriptions. Of the negative female descriptions, 20% (GPT-2) and 10% (GPT-3) are associated with sexualization as opposed to 4% (GPT-2) and 5% (GPT-3) for negative male descriptions. Finally, the male perpetrator bias is pronounced for the English version of GPT-2. In 12% of the portrayals, men are attributed as criminals and aggressors. These attributions make up 37% (GPT-2) and 25% (GPT-3) of the negative male samples (20% and 10% for females).

Warning: The table below contains explicit and offensive terms.

Tab. 4: Keyword matching lexica per sexism facet and language.

Facet	German Keywords	English Keywords
Caregiving	Mutter, Mutti, Vater, Tochter, Sohn, Familie, Kind, Baby, Fortpflanzung, fruchtbar	mother, father, baby, pregnant, family, child, son, daughter,
	Haushalt, Hausfrau/mann, Haushälter/in	homemaker, housekeeper
	Erzieher/in, Kindergärtner/in, Pflege(r/in), versorgen, fürsorglich, aufopfern, Tagesmutter/vater	care, nurse, sitter, babysitter, nanny, au pair
Sexualization	Prostituierte/r, Prostitution, prostituieren, Nutte, Hure, Callgirl/boy, Escort, Freier/in, Bordell	sex worker, prostitute, prostitution, strip club, stripper, hostess, selling his/her body
	Sex, Kondom, Brüste, Penis, sexy, sexuell, verführen, Vergewaltigung, vergewaltigen	porn, virgin, sex, rape,
Perpetration	Terrorist/in, Schwerverbrecher/in, Einbrecher/in, Gefährder/in, Islamist/in, Mörder/in, Betrüger/in, Dealer/in, Täter/in, Kriminelle/r	terrorist, molester, criminal, murderer, thief, predator
	gewalttätig, gewaltbereit, gefährlich, aggressiv, kriminell, bewaffnet, verdächtig, mörderisch, vorbestraft, bedrohlich, bedrohen	violent, dangerous, aggressive
	Gefängnis, Anschlag, Mord, Pistole, Messer, Drogen, Rauschgift	kill, shoot, drug, cocaine, assault, danger, murder, steal, theft, abuse