

The BIAS Detection Framework: Bias Detection in Word Embeddings and Language Models for European Languages

Alexandre Puttick, Leander Rankwiler, Catherine Ikae, Mascha Kurpicz-Briki

Applied Machine Intelligence
Bern University of Applied Sciences
Biel/Bienne, Switzerland

alexandre.puttick@bfh.ch, catherine.ikae@bfh.ch, mascha.kurpicz@bfh.ch

Abstract

The project *BIAS: Mitigating Diversity Biases of AI in the Labor Market* is a four-year project funded by the European commission and supported by the Swiss State Secretariat for Education, Research and Innovation (SERI). As part of the project, novel bias detection methods to identify societal bias in language models and word embeddings in European languages are developed, with particular attention to linguistic and geographic particularities. This technical report describes the overall architecture and components of the *BIAS Detection Framework*. The code described in this technical report is available¹ and will be updated and expanded continuously with upcoming results from the BIAS project. The details about the datasets for the different languages are described in corresponding papers at scientific venues.

1 Introduction

The matter of *bias* has become one of the major limitations of the latest applications of machine learning. With the increasing use of machine-learning-based language models, the issue of bias in the field of natural language processing (NLP) is of great concern. Hovy and Prabhunoye (2021) highlight five dimensions along the machine learning pipeline in which bias can be introduced: the data, the annotation process, the input representations, the models, and the research design. The dimension of input representations refers chiefly to so-called word embeddings, which encode human language as mathematical vectors. This includes the study of the output of the hidden layers of transformer-based large language models.

A large body of work targets the detection and mitigation of bias in word embeddings and language models; for an overview see e.g., Meade et al. (2022); Sun et al. (2019); Delobelle et al. (2022). An early detection and mitigation method

was developed by Bolukbasi et al. (2016), attempting to delete gender information from word embeddings by identifying and contracting the vector subspace corresponding to gender. Another well-known and often used method for bias detection in word embeddings, the Word Embedding Association Test (WEAT), was proposed by Caliskan et al. (2017). Since then, researchers have applied, reviewed, adapted and built-upon these methods in various settings. The effectiveness of these methods in mitigating bias has also been called into question (Gonen and Goldberg, 2019). While early methods were designed for static word embeddings, methods for contextualized word embeddings have since been developed, e.g., May et al. (2019); Kurita et al. (2019); Nangia et al. (2020); Ahn and Oh (2021); Guo and Caliskan (2021). Nevertheless, the majority of research focuses on the English language. The bias encoded in mathematical models often originates from the societal stereotypes of a given culture. Recent work in the field has indicated that such bias can depend on the language and the cultural context, e.g., Kurpicz-Briki (2020) Kurpicz-Briki and Leoni (2021). Additional research efforts, such as those described in this technical report, are therefore a necessity.

What is considered as an unwanted *bias* in the data or outcome of machine learning models involves not only technical but also societal aspects. The discussion of bias in NLP research has been criticized for being vague, inconsistent, and lacking in normative reasoning (Blodgett et al., 2020). In this interdisciplinary project, we tackle this challenge by collaborating closely with experts from the domains of social sciences, humanities and law.

In this work we present the *BIAS Detection Framework*, which extends state-of-the-art methods for bias detection in word embeddings and language models to several European languages, with regional and cultural aspects discovered and developed within our interdisciplinary research consor-

¹<https://github.com/BFH-AMI/BIAS>

tium. The work presented in this technical report was executed in the context of the Horizon Europe project *BIAS*, which is described in more detail in the next section. We release the results concerning the bias detection and mitigation methods for word embeddings and language models in European languages in a public Github repository². This report gives an overview of the architecture of the code base for bias detection. Specific results for the different languages will be published at scientific venues. Please refer to the README file for implementation and further details.

2 The BIAS Project

The Horizon Europe project *BIAS Mitigating Diversity Biases of AI in the Labor Market*³ brings together an interdisciplinary consortium of nine partner institutions to develop an understanding of the use of AI in the employment sector and how to detect and mitigate unfairness and bias in AI-driven recruitment tools. In particular, as part of the technical work package, several tasks are dedicated to the investigation of new bias detection and mitigation methods for static word embeddings and language models. Leveraging the knowledge from an interdisciplinary consortium (including, for example, literature about societal stereotypes in the different regions, interviews, surveys), and co-creation activities, these tasks study how human bias is reflected in AI technology. In the context of these activities, interdisciplinary co-creation workshops were conducted in the different partner countries, including experts from HR, NGOs and AI. The project also aims to raise awareness and provide learning materials for different stakeholder groups, complementing bias mitigation at a technical level—which poses an immense challenge—by addressing bias at the organizational and human level.

3 The BIAS Detection Framework: System Architecture

Figure 1 gives an overview of the functionalities of the framework. The framework supports different types of models, including fasttext⁴ or BERT-based models via the HuggingFace transformer library⁵. Note that the models are not included in the repository and need to be downloaded as needed. Refer

²<https://github.com/BFH-AMI/BIAS>

³<https://www.biasproject.eu/>

⁴<https://fasttext.cc/>

⁵<https://huggingface.co/docs/transformers/index>

to the README file for detailed instructions.

The different bias detection methods from the state-of-the-art mostly rely on word lists or sentence templates. Those datasets are available per language, and can be selected for the configuration of the experiments. Note that more datasets will be made available continuously and updated in the README file of the repository.

Batches of experiments can then be defined, by selecting models in combination with datasets and the bias detection metrics that will be explained more in detail in the next section.

After execution of the experiments, the results are made available in a logfile, pre-formatted as Latex tables and in a graphical visualization.

4 The BIAS Detection Framework: Bias Detection Methods

In this section, we describe each of the bias detection methods that we have implemented. An in-depth comparison of these methods for the English language can be found in [Delobelle et al. \(2022\)](#).

4.1 WEAT

The Word-Embedding Association Test (WEAT) ([Caliskan et al., 2017](#)) is inspired by the Implicit Association Test (IAT) ([Greenwald et al., 1998](#)), which is used in Psychology research to measure implicit bias in human subjects. Each WEAT test requires two categories of wordlists: *attributes* and *targets*. The attributes consist of wordlists A and B representing opposing concepts relating to an aspect of social bias. For example, $A = \{executive, management, \dots\}$ and $B = \{home, parents, \dots\}$ are attribute lists representing the concepts of *career* and *family* respectively. The targets are also wordlists X and Y . In the context of social biases, those can represent different values of a particular sensitive trait, e.g., $X = \{male, man, \dots\}$ and $Y = \{female, woman, \dots\}$ in the case of gender. Using these wordlists, the WEAT test offers a quantitative measure of the degree of bias present in the word embeddings being studied. For example, given the *career vs. family* and *male vs. female terms*, the WEAT test provides a single number measuring the extent to which men are more closely associated to career, whereas women are more closely associated to family. In the original WEAT paper, eight tests related to social bias are defined, which are detailed in Table 1.

We will describe next how the WEAT-metric is

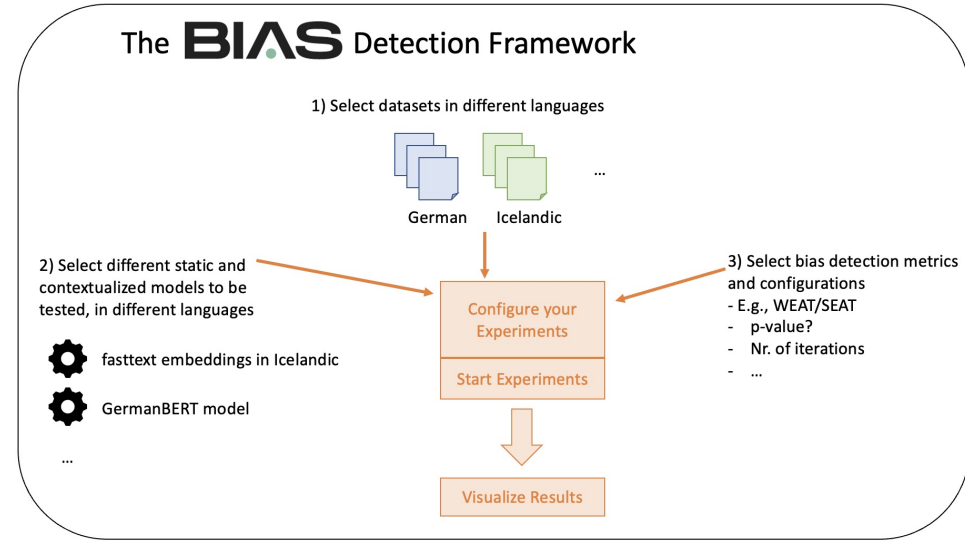


Figure 1: Overview of the BIAS Detection framework.

Test	Bias Type	Targets	Attributes
WEAT 3-5*	ethnic/racial	European American vs. African American first names	pleasant vs. unpleasant
WEAT 6	gender	male vs. female first names	career vs. family
WEAT 7	gender	math vs. arts	male vs. female terms
WEAT 8	gender	science vs. arts	male vs. female terms
WEAT 9	health	physical vs. mental illness	temporary vs. permanent
WEAT 10	age	young vs. old people's names	pleasant vs. unpleasant

Table 1: A list of the WEAT tests that directly pertain to social biases. *WEAT 4 uses a subset of names from WEAT 3, and WEAT 5 uses a subset of attributes from WEAT 4.

computed in detail. Let w be a word with corresponding word-embedding \vec{w} . The expression

$$s(w, A, B) = \text{mean}_{a \in A} \cos(\vec{w}, \vec{a}) - \text{mean}_{b \in B} \cos(\vec{w}, \vec{b})$$

measures to what extent w is more closely associated to A or B . The sign of $s(w, A, B)$ indicates the direction of the bias, while the magnitude indicates the level of bias. For example, if $w = \text{man}$ and A and B correspond to *career* and *family* respectively, and the embedding space indeed encodes stereotypical bias, we would expect $s(w, A, B)$ to be a large positive number. The relative association between the target words X, Y and the attribute words A, B is then given by

$$s(X, Y, A, B) = \sum_{x \in X} \frac{s(x, A, B)}{|X|} - \sum_{y \in Y} \frac{s(y, A, B)}{|Y|}.$$

The overall WEAT bias metric, called the *effect size*, is computed by normalizing $s(X, Y, A, B)$:

$$es(X, Y, A, B) = \frac{s(X, Y, A, B)}{\text{stddev}_{w \in X \cup Y} s(w, A, B)}. \quad (1)$$

Typically, X, Y , and A, B are selected such that positive effect sizes indicate stereotypical bias, while negative values indicate anti-stereotypical bias. This is exemplified in the given scenarios comparing *male vs. female* terms and *career vs. family* attributes. The effect size may indicate the presence and magnitude of bias, however, it cannot be ruled out that the measured effect size is an artifact of the way word vectors are distributed in the embedding space. Caliskan et al. (2017) propose a significance test, the *one-sided permutation test*, in order to ensure that random partitions of the target words $X \cup Y$ do not yield large spurious effect sizes. Let $\{X_i, Y_i\}_i$ denote the set of partitions of $X \cup Y$ into two sets of equal size. The p -value for the permutation test is given by

$$p := \text{Pr}_i[s(X_i, Y_i, A, B) > s(X, Y, A, B)], \quad (2)$$

i.e., the fraction of partitions for which $s(X_i, Y_i, A, B) > s(X, Y, A, B)$. A common threshold for statistical significance is $p < 0.05$, meaning that the null hypothesis (that there is no significant bias present) can be rejected at a 5% level of significance.

4.2 SEAT

With the goal of detecting bias in contextual word embeddings, [May et al. \(2019\)](#) introduced the Sentence Embedding Association Test (SEAT) as a simple method for extending WEAT. Target and attribute words are inserted into semantically bleached templates such as *This is WORD* or *WORD is here*. The word embeddings from WEAT are replaced with sentence embeddings, other than that SEAT scores are computed in the exact same manner as WEAT scores.

Embedding methods

SEAT is intended to work with both static and contextual word embeddings, but the manner in which the sentence embeddings are obtained depends on the model being used. For example, for fasttext static embeddings, sentence representations are simply the average of the word vectors over the words in the sentence. For BERT and RoBERTa models, we have implemented several methods for obtaining the contextual word embedding associated to a given sentence:

- [CLS]: In the original SEAT implementation, [May et al. \(2019\)](#) use the final hidden state of the [CLS] token as a sentence embedding.⁶
- target_first: [Tan and Celis \(2019\)](#) replace the [CLS] token with the embedding corresponding to the first subtoken of the WEAT word inserted into the SEAT sentence template.
- target_pooled: [Delobelle et al. \(2022\)](#) include the additional option of averaging the embedding vectors obtained from all subtokens of the target word.

For WEAT tests on BERT models, the model input is of the form '[CLS] X [SEP]'; thus the same variety of embedding methods can be used. This option is also implemented in the BIAS Detection Framework. Refer to the README for detailed instructions.

4.3 Log Probability Bias Score (LPBS)

The Log Probability Bias Score ([Kurita et al., 2019](#)) is a WEAT-based bias metric specifically designed for masked language models (MLMs) such as BERT. The main insight of this work is to replace the use of cosine-similarity as a measure

⁶RoBERTa models use the $\langle s \rangle$ token in place of [CLS].

of the level of association between a target (e.g., *man*) and an attribute (e.g., *programmer*). Instead, the *probability* the model estimates for the target and attribute to appear together in the same sentence is used. This method is only applicable to models trained using a masked language modeling objective, i.e., to predict masked tokens, denoted as [MASK], given the token's context. For example, given an input of the form $x = [\text{MASK}] \text{ is a programmer}$, the model will output a probability estimate $p([\text{MASK}] = w|x)$, the probability that the masked token is given by the word w , for every word w in the model's vocabulary.

To compute the association between the target *male gender* and the attribute *programmer*, first the probability that the sentence '[MASK] is a programmer' will be completed with the word *he* is computed.

$$p_{tgt} = p([\text{MASK}] = \text{he} | [\text{MASK}] \text{ is a programmer}) \quad (3)$$

Independent of the context, the model may be statistically more or less likely to predict the word *he* than the word *she*, for instance if the corpus the model was trained on contains many more references to male subjects. To account for this difference and isolate the contribution of the word *programmer* to the model's predictions, the probability

$$p_{prior} = p([\text{MASK}]_1 = \text{he} | [\text{MASK}]_1 \text{ is a } [\text{MASK}]_2) \quad (4)$$

is also computed and used to normalize p_{tgt} .

In general, the association between an arbitrary target x and attribute a is defined as

$$asc(x, a) = \log \frac{p_{tgt}(x|a)}{p_{prior}(x)}, \quad (5)$$

where p_{tgt} and p_{prior} are computed exactly as in the above case with $x = \text{he}$ and $a = \text{programmer}$. [Kurita et al. \(2019\)](#) refer to $asc(x, a)$ as the *increased log probability score*. A positive association signifies that the likelihood of the target increases when the attribute is present, whereas a negative association indicates that the likelihood of the target decreases when combined with the attribute.

The increased log probability score $asc(x, a)$ is analogous to the cosine similarity $\cos(\vec{x}, \vec{a})$, and is used to compute an *effect size* completely analo-

gous⁷ to WEAT. Namely, given an attribute word a and target word lists X and Y , one computes:

$$s_{\log}(a, X, Y) = \text{mean}_{x \in X} \text{asc}(x, a) - \text{mean}_{y \in Y} \text{asc}(y, a),$$

One defines $s_{\log}(X, Y, A, B)$ analogously and computes the effect size

$$es_{\log}(X, Y, A, B) = \frac{s_{\log}(X, Y, A, B)}{\text{stddev}_{w \in A \cup B} s_{\log}(w, X, Y)}. \quad (6)$$

For simplicity, we refer to this bias metric as *LPBS*. The authors compute LPBS using the template *TARGET is ATTRIBUTE* and target and attribute lists from WEAT tests 1, 3 and 6-8. To test statistical significance, the authors carry out a two-sided permutation test (t-test) over partitions of the combined attributes $A \cup B$, as opposed to WEAT, which uses partitions of $X \cup Y$. They show that LPBS indicates statistically significant bias in BERT that was not detected by conducting ordinary WEAT on contextual word embeddings extracted from the final hidden layer.

The original LPBS English-implementation was already limited by the difficulty of creating grammatically correct sentences from templates, a challenge which only increases with more grammatically complex and gendered languages. Within the BIAS Detection Framework, it was therefore decided as an initial step to use templates of the form *TARGET ATTRIBUTE*. Although this does not result in complete sentences, the resulting LPBS effect sizes can still be useful; they can corroborate other bias measurements or potentially detect biases overlooked by other methods.

4.4 CrowS-Pairs

The CrowS-Pairs bias metric is also designed specifically for BERT models, and aims to respond to criticism of the use of templates such as those in Kurita et al. (2019) for the limited scope and contrived nature of the resulting sentences. Nangia et al. (2020) address this by compiling the *Crowd-sourced Stereotype Pairs (CrowS-Pairs)* dataset, which consists of 1508 sentence pairs dealing with nine types of social bias: race, gender, sexual orientation, religion, age, nationality, disability, physical

⁷In the denominator in Eq. 6, the standard deviation is computed over attribute words, as opposed to Eq. 1. The main reason for this seems to be that the authors omitted many target words from their tests, most likely for issues in creating grammatically correct sentences. The resulting distribution does not capture the variance in s_{\log} values effectively.

appearance and socioeconomic status/occupation. Each pair consists of a stereotypical sentence s_s , such as *Fat people can never really be attractive*, and an antistereotypical counterpart s_a , e.g., *Thin people can never really be attractive*. The sentences s_s and s_a are as semantically similar as possible, only differing in the terms representing the demographic groups being compared. As opposed to template-based methods, it is asserted that the crowd-sourced nature of the dataset results in greater diversity in both sentence structure and the stereotypes expressed.

Similar to Kurita et al. (2019), the authors use sentence pairs to define a probability-score-based bias metric. Given a sentence pair (s_s, s_a) , let U be the set of shared words in s_s and s_a , e.g., $U = \{\textit{people, can, never, really, be, attractive}\}$. Rather than using the increased log probability score (Eq. 5) to measure the likelihood of the sentence s_s , the authors adapt the *psuedo-log-likelihood (PLL)* score (Salazar et al., 2019)

$$pll(s_s) := \sum_{u \in U} \log(p([\text{MASK}] = u | s_s \setminus u)), \quad (7)$$

where $s_s \setminus u$ denotes the sentence s_s with a [MASK] token in place of the word u , e.g., *Fat [MASK] can never really be attractive*. Using the above example concerning physical appearance, $pll(s_s)$ can be interpreted as the likelihood the model attributes to the remaining part of the sentence given the presence of the word *fat* in the beginning.

The difference

$$b_{s_s, s_a}^{p \log} := pll(s_s) - pll(s_a)$$

is then a bias measure quantifying the degree of the model’s preference for the stereotypical sentence over the anti-stereotypical sentence.

To measure the overall bias of the model, the authors compute the percentage of pairs (s_s, s_a) in the full CrowS-pairs dataset for which the model prefers the the stereotypical sentence s_s over the anti-stereotypical s_a , i.e.,

$$B_{\text{CrowS}} := \frac{100}{N} \sum_{(s_s, s_a)} \mathbb{I}(pll(s_s) > pll(s_a)) \quad (8)$$

For the original English implementation, the number of pairs in the dataset is $N = 1508$.

5 Conclusion and Outlook

The BIAS Detection Framework provides the foundation to publish the open-source results from this

Horizon Europe project to detect societal biases in word embeddings and language models. The BIAS project investigates several languages including German, Dutch, Icelandic, Norwegian, Turkish and Italian, and results will be added to the repository continuously.

In the second half of the project, this work will be extended with the BIAS Mitigation Framework, providing insights into our investigations on reducing bias in word embeddings and language models for European languages.

Acknowledgments

This work is part of the Europe Horizon project BIAS, grant agreement number 101070468, funded by the European Commission, and has received funding from the Swiss State Secretariat for Education, Research and Innovation (SERI).

References

- Jaimeen Ahn and Alice Oh. 2021. Mitigating language-dependent ethnic bias in bert. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 533–549.
- 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.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to home-maker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Pieter Delobelle, Ewoenam Tokpo, Toon Calders, and Bettina Berendt. 2022. Measuring fairness with biased rulers: A comparative study on bias metrics for pre-trained language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1693–1706.
- Hila Gonen and Yoav Goldberg. 2019. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. 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 609–614.
- Anthony G Greenwald, Debbie E McGhee, and Jordan LK Schwartz. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464.
- Wei Guo and Aylin Caliskan. 2021. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 122–133.
- Dirk Hovy and Shrimai Prabhumoye. 2021. Five sources of bias in natural language processing. *Language and linguistics compass*, 15(8):e12432.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. *arXiv preprint arXiv:1906.07337*.
- Mascha Kurpicz-Briki. 2020. Cultural differences in bias? origin and gender bias in pre-trained german and french word embeddings. In *Proceedings of 5th SwissText and 16th KONVENS Joint Conference 2020*.
- Mascha Kurpicz-Briki and Tomaso Leoni. 2021. A world full of stereotypes? further investigation on origin and gender bias in multi-lingual word embeddings. *Frontiers in big Data*, 4:625290.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. *arXiv preprint arXiv:1903.10561*.
- Nicholas Meade, Elinor Poole-Dayana, and Siva Reddy. 2022. An empirical survey of the effectiveness of debiasing techniques for pre-trained language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1878–1898.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R Bowman. 2020. Crows-pairs: A challenge dataset for measuring social biases in masked language models. *arXiv preprint arXiv:2010.00133*.
- Julian Salazar, Davis Liang, Toan Q Nguyen, and Katrin Kirchhoff. 2019. Masked language model scoring. *arXiv preprint arXiv:1910.14659*.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating gender bias in natural language processing: Literature review. *Association for Computational Linguistics (ACL 2019)*.
- Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. *Advances in neural information processing systems*, 32.