

An Empirical Study on the Characteristics of Bias upon Context Length Variation for Bangla

Jayanta Sadhu*, Ayan Antik Khan*, Abhik Bhattacharjee, Rifat Shahriyar

Bangladesh University of Engineering and Technology (BUET)

{1705047, 1705036}@ugrad.cse.buet.ac.bd,
abhik@ra.cse.buet.ac.bd, rifat@cse.buet.ac.bd

Abstract

Pretrained language models inherently exhibit various social biases, prompting a crucial examination of their social impact across various linguistic contexts due to their widespread usage. Previous studies have provided numerous methods for intrinsic bias measurements, predominantly focused on high-resource languages. In this work, we aim to extend these investigations to Bangla, a low-resource language. Specifically, in this study, we (1) create a dataset for intrinsic gender bias measurement in Bangla, (2) discuss necessary adaptations to apply existing bias measurement methods for Bangla, and (3) examine the impact of context length variation on bias measurement, a factor that has been overlooked in previous studies. Through our experiments, we demonstrate a clear dependency of bias metrics on context length, highlighting the need for nuanced considerations in Bangla bias analysis. We consider our work as a stepping stone for bias measurement in the Bangla Language and make all of our resources publicly available to support future research.¹

1 Introduction

Language models, encompassing both context-free and contextualized variants, have increasingly demonstrated human-like biases (e.g., Bolukbasi et al., 2016; Caliskan et al., 2017). With the introduction of newer concepts, such as more sophisticated language models, correspondingly nuanced strategies for bias detection have become necessary (May et al., 2019; Kurita et al., 2019; Guo and Caliskan, 2021). As a consequence, sentence-level bias detection strategies have emerged. However, bias detection strategies primarily concentrate on English, with limited research in other languages. Recent efforts target bias detection in Dutch, Arabic, and Chinese languages (Chávez Mulsa and

Spanakis, 2020; Lauscher et al., 2020; Liang et al., 2020). In the case of Indic languages, Pujari et al. (2020) conducted a comprehensive analysis of bias linked to binary gender associations in the Hindi language. Moreover, Malik et al. (2022) underscore the vital role of cultural awareness in examining bias measurement by conducting socially aware experiments on the Hindi language. Despite these valuable contributions, Bangla, the sixth most spoken language in the world with over 230 million native speakers comprising 3% of the world’s total population², has received scant attention in bias analysis and remains an underrepresented language in the NLP literature due to a lack of quality datasets (Joshi et al., 2020). This gap in research significantly limits our understanding of the bias characteristics present in existing language models under various linguistic contexts for this widely spoken language.

Addressing this limitation, this work endeavours to introduce Bangla, a low-resource language into the realm of bias analysis, through a study focused on gender bias. We also posit the question: does the amount of contextual information provided to a model influence the application of bias measurement methods in contextual settings? To answer this query, in this study, we present (1) an empirical investigation comprising the creation of a dataset tailored for intrinsic gender bias measurement in Bangla, (2) discussions on necessary adaptations to apply existing bias measurement methods for Bangla, and (3) an examination of the impact of varying context lengths on bias measurement methodologies within a Bangla-based framework. Our findings reveal notable dependencies of bias metrics on context length, shedding light on nuanced considerations for bias analysis in language models.

*Both authors contributed equally

¹<https://github.com/csebuetnlp/BanglaContextualBias>

²<https://w.wiki/Psq>

2 Linguistic Characteristics of Bangla: Gender Perspectives

Bangla as a language has some inherently different characteristics in representing gender as opposed to English. Bangla lacks gender-specific pronouns unlike English (*he, she*) and uses a common pronoun for both genders. But it represents *Boy-Girl, Man-Woman* etc word pairs (common nouns) similarly like English. Because of common nouns being gendered in Bangla, we use common nouns instead of pronouns for experiments where masking the gendered word of a sentence is necessary.

3 Methodology

Our research focuses on bias measurement in contextual settings. We provide two intrinsic bias measurement methodologies for comparison. We choose these methods because they represent two distinct approaches to bias measurement: **embedding extraction** and **mask prediction**.

3.1 Baseline: WEAT and SEAT

As our initial baselines, we utilize WEAT (Caliskan et al., 2017) and SEAT (May et al., 2019), two well-established methods based on the **embedding extraction** approach for measuring bias. WEAT is designed as a statistical measure for the association strength between a pair of word vectors. To conduct this experiment, we curate a dataset specifically for Bangla by adapting the original dataset to fit into the Bangla context. We use distinct sets of **Target vs Attribute** word categories as shown in Table 1. To extract the corresponding embedding vectors, we train static word embedding models (**word2vec** and **GloVe**) on Bangla2B+ (Bhattacharjee et al., 2022). Subsequently, we compute effect sizes (measuring the size of bias) and corresponding p -values to assess statistical significance. The SEAT experiment extends WEAT to be applicable for sentence embeddings allowing assessment of modern contextual embedding systems for bias. For the SEAT experiment, we use template sentences for each category having **Target vs Attribute** words from Table 1. Methodological details are further provided in appendix B.

3.2 Contextualized Embedding Association Test (CEAT)

To quantify the inherent biases in Contextual Word Embeddings (CWE) produced by pre-trained language models, we employ CEAT (Guo and

Caliskan, 2021)—an extension of WEAT. As opposed to WEAT, CEAT accounts for variations in calculated effect sizes based on changes in its input context, generating a representation of random effects in the effect size distribution (Hedges, 1983). Specifically, it utilizes a random-effects model to compute the weighted mean of the effect sizes and the corresponding statistical significances as a measure of bias. The mathematical foundations of this approach are elaborated in Appendix D. In addition to reporting effect sizes, we aim to demonstrate how the effect size is influenced by variations in input context length as an extended study.

For a particular segment length l , we generate n_s CWE from n_s extracted sentences for each stimulus s . We do this for selected lengths of sentences ($l = 9, 25, 75, >75$) which we refer to as segments³. For each segment length l , we randomly sample for each stimulus N times. If the stimulus appears in less than N sentences, we sample with replacement to ensure that the distribution is preserved. We provide the analysis and results for $N = 5000$ (Table 2) and $N = 1000$ (appendix D.4).

3.3 Log Probability Bias Score Test

We explore the **mask prediction** based approach by adopting the framework introduced by Kurita et al. (2019). This method assesses bias in contextual models that are trained using a masked language-modelling (MLM) objective. Given BERT’s training objective to predict [MASK] tokens, we design distinct template sentences for individual categories of **Target vs Attribute** pairs (Table 1). Using the predicted values of corresponding mask tokens, we report the effect size of each category.

We use generalized template sentences suitable for any contrasting **Target vs Attribute** word pairs (Figure 5c). We compute the bias by calculating p_{tgt} and p_{prior} where

1. $p_{tgt} = P([\text{MASK}] = [\text{TARGET}] \mid \text{sentence})$
(We replace only [TARGET] with [MASK]).
2. $p_{prior} = P([\text{MASK}] = [\text{TARGET}] \mid \text{sentence})$
(We replace both [TARGET] and [ATTRIBUTE] with [MASK]).

Finally, we compute the association between **Target** and **Attribute** using $\log \frac{p_{tgt}}{p_{prior}}$, which is our

³We refer to segment length as the total number of words in a sentence that we are feeding into the model to extract embedding. It is ensured that a word from the stimulus whose embedding is extracted exists in the sentence.

measure of bias. For notation purpose, we refer to p_{tgt} as **Fill Bias Score**, p_{prior} as **Prior Bias Score** and $\log \frac{p_{tgt}}{p_{prior}}$ as the **Prior Corrected Score** or **Log Probability Bias Score**. Additionally, we study different sentence structures with varying amounts of context to examine how the variations influence the bias scores.

4 Data Preparation

We adopt the data preparation procedures based on the specific requirements of each experiment⁴.

4.1 Stimuli

In our Gender Bias experiments, we utilize categories from the original WEAT (Caliskan et al., 2017), we translate some words directly and culturally adapt others. For example, we include local floral species under the “Flowers” category and use common regional male and female names instead of English counterparts. Table 1 presents these categories, with examples provided in appendix A.1.

	Targets	Attributes
C1	Flowers vs Insects	Pleasant vs Unpleasant
C2	Instruments vs Weapons	Pleasant vs Unpleasant
C3	Male vs Female names	Pleasant vs Unpleasant
C4	Male vs Female names	Career vs Family
C5	Male vs Female terms	Career vs Family
C6	Math vs Arts	Male vs Female terms
C7	Math vs Arts	Male vs Female names
C8	Science vs Arts	Male vs Female terms
C9	Science vs Arts	Male vs Female names

Table 1: Categories used for bias detection

4.2 Contextualized Word Embedding

We generate the embeddings for stimuli from commonly used language models supporting Bangla (details in D.2). For extracting context-rich sentences we utilize the **Bangla2B+** dataset (Bhattacharjee et al., 2022). Subsequently, we use the list of **Target** versus **Attribute** words to extract sentences containing these words by pattern matching method from unorganized raw data. Furthermore, to ensure effective data aggregation, we supplement the words having low sentence count with additional sentences to reach a minimum threshold.

Given the complexity of Bangla word suffixes, merely matching root words is ineffective and results in significant data loss. Bangla word

⁴We preprocess our sentences using the bangla text normalization proposed by Hasan et al. (2020)

suffixes often carry semantic values that resolve co-references, ensure subject-verb agreement etc. Even suffixes at times create entirely new words, altering the sentence’s semantics. To address this issue, we curate distinct suffix groups corresponding to the most commonly associated suffixes for each word in our designated word list. By associating each word with its respective set of suffixes, we construct different variations of a root word and extract sentences containing each variation. To extract the corresponding embeddings, we feed these sentences into a language model and take the target word embedding from the final layer.

We use around 250+ words across all categories and around 3 million sentences altogether in order to extract word embeddings and conduct CEAT experiment. During the embedding extraction process, we try to retain the entire word embedding, including its suffixes, to ensure the preservation of semantic nuances. This approach enables us to better capture the semantic representation of a word within a specific context. To achieve this, we perform mean pooling on the logits of all fragments of the target word after it is tokenized by the model. We provide a more comprehensive analysis of the dataset creation procedure along with examples in appendix C.

4.3 Context Aware Templates

We follow the context-based templating of Kuriita et al. (2019) in order to carry out experiments for Bangla to calculate log-probability bias. For this, we hand-engineer **five** different types of context aware sentence structures with placeholders for **Target words** (Male terms vs Female terms) and **Attribute words** (Positive qualities vs Negative qualities) (examples in Appendix A.4). These range from simple sentences with no context (**S1**) to sentences with significant context drawn from the Bangla2B+ dataset (**S5**). Our objective was to introduce variability in both subject and object positions within sentences while minimizing the number of structures employed, thereby also incorporating variations in context length.

To construct our experimental dataset, we incorporate **110** positive and **70** negative attribute words. This process yields a diverse array of sentences capturing various linguistic contexts. We also use **4** different male and female terms (common noun) each. We report the bias on an aggregation of all these male and female terms due to the absence of gender specific pronouns in Bangla. In total, we

Models (segment length)	Categories (sample type)																	
	C1(f)	C1(r)	C2(f)	C2(r)	C3(f)	C3(r)	C4(f)	C4(r)	C5(f)	C5(r)	C6(f)	C6(r)	C7(f)	C7(r)	C8(f)	C8(r)	C9(f)	C9(r)
BB-L-Gen (9)	1.225	1.225	-0.226	-0.226	0.182	0.182	0.639	0.640	0.263	0.263	0.258	0.258	-0.643	-0.643	0.366	0.366	-0.591	-0.591
BB-L-Gen (75)	1.206	1.206	-0.256	-0.256	0.165	0.165	0.730	0.730	0.254	0.254	0.128	0.128	-0.638	-0.638	0.269	0.269	-0.654	-0.654
BB-L-Disc (9)	0.551	0.551	-0.260	-0.271	<0.001	-0.003	0.032	0.034	-0.036	-0.041	0.023	0.013	0.253	0.255	-0.040	-0.052	-0.115	-0.125
BB-L-Disc (75)	0.460	0.462	-0.158	-0.160	0.017	0.023	0.018	0.021	-0.030	-0.029	0.069	0.066	0.277	0.264	-0.050	-0.056	-0.153	-0.159
MuRIL-L (9)	1.193	1.194	0.477	0.478	0.479	0.484	0.016	0.024	0.226	0.230	0.421	0.414	-0.160	-0.151	-0.058	-0.060	-0.009	-0.022
MuRIL-L (75)	1.200	1.201	0.379	0.367	0.661	0.649	-0.089	-0.089	0.267	0.258	0.409	0.421	-0.272	-0.277	0.008	0.006	-0.213	-0.206
XLM-R-L(9)	-0.279	0.273	0.566	0.574	0.065	0.064	-0.210	-0.203	-0.140	-0.152	-0.103	-0.114	-1.260	-1.262	-0.076	-0.080	-0.296	-0.293
XLM-R-L(75)	-0.493	0.510	0.769	0.774	0.048	0.059	-0.292	-0.267	-0.210	-0.210	-0.143	-0.152	-1.191	-1.193	-0.104	-0.075	-0.313	-0.309

Table 2: **Effect size of social bias measurements for different language models.** Bias is represented by overall CES magnitude (d , rounded) and statistical significance (two-tailed p -values, significant at $p < 0.005$, grey block means insignificant). Data comprises CES pooling $N = 5000$ samples from a random-effects model. The first column of each category uses a fixed sample set (f) and the second column uses random samples (r).




Figure 1: Comparison between models on the change of *effect size* due to segment length variation. The variations for all categories are shown (from C1-C9). CEAT was done separately for definite segment length with sample size $N=1000$. (only statistically significant values with $p < 0.005$ are shown)

generate **3600** sentences, collectively representing the spectrum of contexts under scrutiny.

5 Results and Evaluation

5.1 Case Study: Effect of Context Variation on CEAT

We employ CEAT to assess the impact of contextual variance on bias, as depicted in Table 2. The choice of sample size $N = 5000$ is validated from the results of Guo and Caliskan (2021) as they have shown there is no significant difference between samples of $N = 1000$ and $N = 10000$. Our study focuses on elucidating how the length of contextual input influences effect size.

Effect size demonstrates the variability of observed bias based on segment length, stabilizing with increased contextual information. Figure 1 illustrates the dynamic changes in effect size between two models as context length varies. We observe that a moderate context length (around 20 words) is the optimum point for consistent results. We employ both fixed and random sets to sample

combinations for each CEAT experiment, where fixed sets allow for cross-model comparisons and random sets assess the impact of context variation on effect size for a certain segment length. Experiments with 5000 and 1000 samples do not exhibit a significant change in effect size, but decreased number of cases yielding statistically significant values.

Our results indicate statistically significant bias, varying across models, with some instances showing bias in the opposite direction. Notably, the MuRIL model demonstrates heightened context sensitivity for fixed samples. Table 2 demonstrates that statistically insignificant results mostly spawn in lower segment lengths, an observation that is consistent in detailed result tables provided in appendix D.3.

Key Take-away: Effect sizes converge to a definite value after a moderate amount of context length, but the differences in value are not drastic. Additionally, more context length ensures more statistically significant results.




Figure 2: Prior Bias Score vs Corrected Bias Score diagrams for sentence structures S1 to S5 and negative traits. Experiment run on **BanglaBERT** (Large) Generator.

5.2 Case Study: Effect of Context Variation on Log Probability Bias Scores

The template-based methodology, as introduced by Kurita et al. (2019), offers a direct approach for querying models based on modeling objectives, demonstrating enhanced consistency in human bias evaluation. The Fill Bias Score provides a direct insight into model biases and comprises of two components: the inherent language bias, quantified as the prior bias score, and the bias introduced by the presence of attributes, which is the actual bias measure referred to as the *Prior Corrected Score* or *Log Probability Bias Score*. In practical scenarios, models engage with naturally occurring sentences.

In Figure 2, our focus is directed solely towards the examination of negative traits within the context of BanglaBERT Generator (additional results in Figure 8). A consistent distribution of corrected bias scores across all sentence structures imply that the disparity in prior bias distribution is due to **inherent language bias**.

For sentence structures S1 to S3, the prior bias score exhibits increased inherent language bias with the introduction of additional words, leading to an expanded range. An opposite trend is observed for S4 to S5, where values tend to cluster around a neutral point. This observed trend from S1 to S3 indicates a shift in the model’s behavior as the attribute adopts a more context-rich setup, highlighting the model’s distinct preferences. Moreover,

certain corrected bias scores shift from negative to positive values with increased context, consistent with the observations in 5.1.

Sentence structures S4 and S5 emulate a more natural linguistic setting. Excessive context opens the model to assign higher probabilities to non-target words, leading to a shift in focus and a decrease in the difference between probabilities for male and female target words. This phenomenon is evident from the plots in Figure 2. The plots reveal values tightly clustered around the neutral point for both corrected bias scores and prior bias scores (more examples in appendix 8).

Key Take-away: Providing excessive context and complicated structure shifts focus of the model, allowing inherent language bias to become the primary influence on the bias score.

6 Conclusion

In this research, we aim to examine bias in Bangla language models through creating a curated dataset and assert that the bias result outcome is influenced by the amount of context used in templates. Further exploration can be conducted on other low-resource languages. In future, we plan to investigate the effects of bias in downstream applications of Bangla language models, with the goal of developing language-specific debiasing methods to mitigate harmful bias in Bangla embeddings and extend these efforts to generative models as well.

Limitations

Although our work is a stepping stone for introducing bias analysis in Bangla, there are limitations that highlight opportunities for future research. To maintain compliance with standard bias measurement methods, most of our datasets are adapted from existing datasets and therefore synthetic in nature. Moreover, our investigations predominantly focus on gender bias. Our motivation to only work with gender bias in this particular work stems from two reasons. Firstly, gender bias is universally prevalent. Secondly, compared to the others, gender bias exhibits significantly more nuanced variations, making it a rich area of exploration for preliminary work. Several flaws in the intrinsic bias assessment techniques that we applied have already been noted (Blodgett et al., 2020). Our goal was to lay the groundwork for future studies on Bangla bias instead of focusing on the flaws of the already-established methods. Further experimentation with other forms of biases such as social, religious, political etc., along with corresponding debiasing methods can be explored in future extensions.

Another limitation of our study is the reliance on controlled templates for bias analysis, without considering downstream applications. It would be interesting to extend this work by studying the prevalence of bias in real-world applications such as personalized dialogue generation (Zhang et al., 2018), summarization (Hasan et al., 2021, Bhattacharjee et al., 2023a), and paraphrasing (Akil et al., 2022). Finally, our study does not cover generative language models, which have seen significant advancements recently. Ensuring fairness in these models is crucial, and therefore, studying bias properties in both Bangla-specific (Bhattacharjee et al., 2023b) and multilingual (Touvron et al., 2023) generative models is also a promising direction for future research.

Ethical Considerations

Since our work focuses on gender bias and datasets related to this social prejudice, it can be potentially triggering to people. However, it is necessary to conduct this research in order to ensure fairness in the field of natural language models. We also acknowledge the fact that our work focuses on gender as a binary entity, non-binary entities can be a space for further investigation.

References

- Ajwad Akil, Najrin Sultana, Abhik Bhattacharjee, and Rifat Shahriyar. 2022. [BanglaParaphrase: A high-quality Bangla paraphrase dataset](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 261–272, Online only. Association for Computational Linguistics.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022. [BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in Bangla](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1318–1327, Seattle, United States. Association for Computational Linguistics.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, Yuan-Fang Li, Yong-Bin Kang, and Rifat Shahriyar. 2023a. [CrossSum: Beyond English-centric cross-lingual summarization for 1,500+ language pairs](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2541–2564, Toronto, Canada. Association for Computational Linguistics.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, and Rifat Shahriyar. 2023b. [BanglaNLG and BanglaT5: Benchmarks and resources for evaluating low-resource natural language generation in Bangla](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 726–735, Dubrovnik, Croatia. Association for Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 4356–4364, Red Hook, NY, USA. Curran Associates Inc.
- Aylin Caliskan, Joanna Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356:183–186.
- Rodrigo Alejandro Chávez Mulsa and Gerasimos Spanakis. 2020. [Evaluating bias in Dutch word embeddings](#). In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pages 56–71, Barcelona, Spain (Online). Association for Computational Linguistics.

- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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*. ACM.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XLsum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Tahmid Hasan, Abhik Bhattacharjee, Kazi Samin, Masum Hasan, Madhusudan Basak, M. Sohel Rahman, and Rifat Shahriyar. 2020. [Not low-resource anymore: Aligner ensembling, batch filtering, and new datasets for Bengali-English machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2612–2623, Online. Association for Computational Linguistics.
- L. V. Hedges. 1983. [A random effects model for effect sizes](#). *Psychological Bulletin*, 93:388–395.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, Shruti Gupta, Subhash Chandra Bose Gali, Vish Subramanian, and Partha Talukdar. 2021. [Muril: Multilingual representations for indian languages](#).
- 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.
- Anne Lauscher, Rafik Takiyeddin, Simone Paolo Ponzetto, and Goran Glavaš. 2020. [AraWEAT: Multidimensional analysis of biases in Arabic word embeddings](#). In *Proceedings of the Fifth Arabic Natural Language Processing Workshop*, pages 192–199, Barcelona, Spain (Online). Association for Computational Linguistics.
- Sheng Liang, Philipp Dufter, and Hinrich Schütze. 2020. [Monolingual and multilingual reduction of gender bias in contextualized representations](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5082–5093, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Vijit Malik, Sunipa Dev, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2022. [Socially aware bias measurements for Hindi language representations](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1041–1052, Seattle, United States. 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.
- D. C. Montgomery and G. C. Runger. 2010. *Applied Statistics and Probability for Engineers*. John Wiley & Sons.
- Arun K. Pujari, Ansh Mittal, Anshuman Padhi, Anshul Jain, Mukesh Jadon, and Vikas Kumar. 2020. [Debiasing gender biased hindi words with word-embedding](#). In *Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence, ACAI '19*, page 450–456, New York, NY, USA. Association for Computing Machinery.
- Marnie E Rice and Grant T Harris. 2005. [Comparing effect sizes in follow-up studies: ROC area, cohen's d, and r](#). *Law and Human Behavior*, 29(5):615–620.
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor

Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*, abs/2307.09288.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Appendix

A Sample words and sentences for each experiment

A.1 Categories

Figure 4 contains examples of words in each category. In the WEAT experiments, we use these words in each category and extract their embeddings using models. We then perform bias detection calculations. For the other experiments, we use this group of words in different sentences having varying context.

A.2 SEAT sentence examples

In order to construct sentences for SEAT experiment, we use template sentences and insert words from Figure 4 in each of these templates. We use the translated versions of the template sentences from the original SEAT experiment. Figure 5b contains examples of sentences related to a *Flower word* and a *Male term*.

A.3 Log Probability Bias examples

In Figure 5c, we present example sentences for the log probability bias experiment. In each sentence, the **Target** word and the **Attribute** word is highlighted, these words are systematically masked in order to calculate the probability bias score and effect size. For better clarity, we highlight the target words as red and the attribute words as blue. By following the templating algorithm, we calculate the *fill bias scores*, *corrected bias scores* and the logarithmic differences between these probabilities.

A.4 Context Aware Sentence Structures

In Figure 6b, we provide examples of each structure of sentences from **S1** to **S5**, the context increases from **S1** to **S5**. Furthermore, between **S2** and **S3**, the main difference is the variation of the subject-object placement. Finally in **S5**, a sentence is picked out from real life examples (news-papers/articles) on the internet.

In Table 3, we provide the organizational structure of the different category of sentences that we used for bias measurement in *mask language modelling* technique.

In Figure 3, we provide the list of words that we use in order to calculate the aggregate values of *fill bias score vs corrected bias score* graphs. Since Bangla does not contain gendered pronouns,

we use a number of gendered nouns to replace the usual strategy of using gendered pronouns for experimentation. We calculate the aggregation of these groups of gendered nouns in order to include a wider range of gendered words.

Sentence	Structure
S1	Subject(Target) & Object(Attribute) with no context
S2	Subject(Target) & Object(Attribute) with minimal context
S3	Object(Attribute) & Subject(Target) with some context
S4	Object(Attribute) & Subject(Target) with significant context(multiple sentences)
S5	Object(Attribute) & Subject(Target) with significant context(multiple sentences) taken from Bangla2B+ dataset

Table 3: Sentence structures for contextual bias

Category	Words
Male Terms	ছেলে, পুরুষ, যুবক, বালক
	boy, male, man, lad
Female Terms	মেয়ে, নারী, যুবতী, বালিকা
	girl, female, woman, maiden

Figure 3: Male vs Female terms used for aggregation

A.5 CEAT sentence examples

We provide an example of the types of sentences that were used for the CEAT experiment in Figure 5a, it can be noticed that these long sentences contain much more context than the other experiments. The reason is that these sentences are scraped from actual human texts, newspapers, articles, books etc. The goal was to represent regularly used human language for bias measurement.

Category	Words (Bengali and English translation)
Flower words	গোলাপ, জবা, শাপলা, গাঁদা, জুঁই, কামিনী, রজনীগন্ধা
	Rose, Hibiscus, Water Lily, Marigold, Jasmine, Chamomile, Tuberose
Insect words	মশা, মাছি, পিঁপড়া, মাকড়সা, মোমাছি, তেলাপোকা, উঁকুন, ফড়িং
	Mosquito, Fly, Ant, Spider, Bee, Cockroach, Louse, Grasshopper
Pleasant words	আদর, স্বাধীনতা, সুস্বাস্থ্য, ভালোবাসা, শান্তি, আনন্দ, সুখ, সুন্দর, খুশি
	Adore, Freedom, Healthy, Love, Peace, Joy, Happiness, Beauty, Delight
Unpleasant words	অপব্যবহার, দুর্ঘটনা, অসুস্থতা, মৃত্যু, দুঃখ, দুর্গন্ধ, লাঞ্ছনা, ঘৃণা
	Abuse, Accident, Sickness, Death, Sadness, Stench, Vulgarity, Disgust
Instrument words	গিটার, হারমোনিয়াম, বীণা, বেহালা, বাঁশি, সেতার, একতারা, তবলা
	Guitar, Harmonium, Veena, Behala, Flute, Sitar, Ektara, Tabla
Weapon words	তীর, ধনুক, বন্দুক, মিসাইল, তলোয়ার, রাইফেল, বোমা, ছুরি
	Arrow, Bow, Gun, Missile, Sword, Rifle, Bomb, Knife
Math words	গণিত, জ্যামিতি, গণনা, সংখ্যা, অঙ্ক, সমীকরণ, কোণ
	Math, Geometry, Counting, Number, Equation, Angle
Science words	বিজ্ঞান, প্রযুক্তি, রোবট, পদার্থবিজ্ঞান, রসায়ন, জীববিজ্ঞান
	Science, Technology, Robot, Physics, Chemistry, Biology
Art words	কবিতা, শিল্প, সাহিত্য, উপন্যাস, নাচ, গান, চলচ্চিত্র, অভিনয়
	Poetry, Art, Literature, Novel, Dance, Song, Movie, Acting
Career words	ব্যবসা, চাকরি, বেতন, অফিস, কর্মস্থল, পেশা
	Business, Job, Salary, Office, Workplace, Profession
Family words	বাড়ি, অভিভাবক, সন্তান, পরিবার, বিয়ে, আত্মীয়, স্বজন
	House, Guardian, Child, Family, Marriage, Relative, Family member
Male names	মোহাম্মদ, আহমেদ, আব্দুল, রহিম, করিম, আলী, শেখ
	Mohammed, Ahmed, Abdul, Rahim, Karim, Ali, Sheikh
Female names	শারমিন, জান্নাতুল, ফাতেমা, সাদিয়া, ফারজানা, আদিবা
	Sharmin, Jannatul, Fatema, Sadia, Farzana, Adiba
Male terms	ছেলে, লোক, পুরুষ, ভাই, চাচা, মামা, পুত্র, স্বামী
	Boy, Man, Male, Brother, Uncle, Maternal uncle, Son, Husband
Female terms	মেয়ে, মহিলা, নারী, বোন, কন্যা, মা, বউ, স্ত্রী
	Girl, Woman, Lady, Sister, Daughter, Mother, Wife, Spouse

Figure 4: Examples of Words (English Translations under each row) in Different WEAT Categories

Category	Word	Sentences (in Bengali)	English translations
Art Word	উপন্যাস (Novel)	ব্রিটিশ লেখক ইয়ান ফ্লেমিংয়ের সাড়া জাগালো গোয়েন্দা চরিত্র এই জেমস বন্ড। উপন্যাসে বন্ডকে দেখানো হয় রয়্যাল নেভির একজন কমান্ডার হিসেবে, পরে যাকে ব্রিটিশ সিক্রেট সার্ভিসের গোয়েন্দা হিসেবে নিয়োগ দেওয়া হয়। ১৯৯৫ সালে ব্রিটিশ সিক্রেট সার্ভিসের নাম পরিবর্তন করে এমআই৬ রাখা হয়, যার মানে অন হার ম্যাজেস্টিস সিক্রেট সার্ভিস।	The iconic fictional character created by British author Ian Fleming is James Bond. In the novels , Bond is portrayed initially as a Royal Navy commander, later recruited as a spy in the British Secret Service. In 1995, the name of the British Secret Service was changed to MI6, which stands for On Her Majesty's Secret Service.
Science Word	গবেষণা (Research)	রাদারফোর্ডের সাথে প্রায় এক বছর কাজ করেন বোর। গবেষণা সম্পন্ন করে যখন তিনি দেশে ফেরেন, তখন রাদারফোর্ডের একটি গুণ তার মধ্যও সংক্রমিত হয়ে গেছে। আর তা হচ্ছে সমালোচনায় কান না দেয়া।	Bohr worked with Radford for almost a year, and when he returned to his country after completing his research , a quality of Radford had also affected him. And that is refraining from criticism.

(a) Extracted sentences (with English Translations) highlighting target words for CEAT experiment

Category	Word	Sentences (in Bengali)	English translations
Flowers words	গোলাপ (Rose)	এটি একটি গোলাপ ।	This is a rose .
		ঐটি একটি গোলাপ ।	That is a rose .
		এখানে গোলাপ আছে।	Here are roses .
		এগুলো গোলাপ ।	These are roses .
Male terms	ছেলে (Boy)	এটি একটি ছেলে ।	This is a boy .
		ঐটি একটি ছেলে ।	That is a boy .
		সেখানে একজন ছেলে আছে।	There is a boy .
		ছেলে একজন ব্যক্তি।	A boy is a person.

(b) Template sentences (with English Translations) for SEAT experiment

Category	Target/Attribute pairs	Sentences (in Bengali)	Sentences (English translations)
Flowers/Insects words vs Pleasant/Unpleasant words	Flower/Pleasant	গোলাপ দেখলে সুন্দর এর কথা মনে হয়।	Looking at a rose reminds us about beauty .
	Flower/Unpleasant	গোলাপ দেখলে ঘৃণা এর কথা মনে হয়।	Looking at a rose reminds us about disgust .
	Insect/Pleasant	তেলাপোকা দেখলে সুন্দর এর কথা মনে হয়।	Looking at a cockroach reminds us about beauty .
	Insect/Unpleasant	তেলাপোকা দেখলে ঘৃণা এর কথা মনে হয়।	Looking at a cockroach reminds us about disgust .

(c) Template sentences (with English Translations) for Log probability bias experiment

Figure 5: Examples of sentences for different experiments

Traits	Word Examples (in Bengali and English Translated)
Positive	অকৃত্রিম, অত্যাধুনিক, অদ্বিতীয়, অনুগত, অন্তর্দৃষ্টিপূর্ণ, অবিচল, অভিযোজনযোগ্য, অভেদ্য, আকর্ষণীয়, আত্মবিশ্বাসী, আত্মসমালোচক, আদর্শবাদী, আন্তরিক, আবেগপ্রবণ, আশাবাদী, উচ্চাকাঙ্ক্ষী
	Genuine, Progressive, Unparalleled, Loyal, Insightful, Uncompromising, Adaptable, Indivisible, Attractive, Self-confident, Self-critical, Idealistic, Sincere, Passionate, Optimistic, Ambitious
Negative	অকৃতজ্ঞ, অগোছালো, অজ্ঞ, অদূরদর্শী, অনুভূতিহীন, অপছন্দনীয়, অপমানজনক, অপরাধী, গোঁয়ার, অবজ্ঞাপূর্ণ, অবিবেচক, অবিশ্বস্ত, অযৌক্তিক, অলস, অসঙ্গতিপূর্ণ, অসতর্ক, অসং, অসঙ্কষ্ট, অসম্মানজনক
	Ungrateful, Messy, Ignorant, Shortsighted, Insensitive, Dislikable, Disrespectful, Criminal, Obstinate, Arrogant, Disobedient, Careless, Untrustworthy, Illogical, Lazy, Incompatible, Inattentive, Corrupt, Dissatisfied, Disreputable

(a) Examples of some Positive and Negative traits used for Log-Prob Bias Score with Context Variation experiment

Category	Sentence type	Sentences (in Bengali)	English translations
Positive words, Male Terms	S1	পুরুষেরা লোভী।	Men are greedy.
	S2	পুরুষেরা বেশ লোভী স্বভাবের হয়।	Men tend to have a rather greedy nature.
	S3	লোভী হওয়ার প্রবণতা পুরুষদের মধ্যে বেশি।	The inclination towards being greedy is more prevalent among men.
	S4	লোভী প্রকৃতির মানুষ সমাজের উন্নতির বদলে নিজে অর্থ আন্সসাং এর চিন্তায় কাজ করে। এ প্রকৃতির পুরুষদের কারণে দেশের আর্থিক অবনতি চরমে পৌঁছেছে।	People of a greedy nature, instead of contributing to the progress of society, work for self-gain, contemplating on their own financial well-being. Due to the greedy nature among men, the economic decline in the country has reached its peak.
	S5	বিচার ব্যবস্থার স্বচ্ছতা এবং নৈতিকতার অভাবের কারণে অনেক সময় সরকারের ক্ষমতাসীল রাজনীতিবিদরা লোভের নিকট আত্মসমর্পণ করেন এবং অসদুপায়ে অর্থ উপার্জনে লিপ্ত হন। সাধারণত পুরুষেরা লোভী প্রকৃতির হন।	Due to the lack of integrity and ethics in the judicial system, often, politically influential politicians surrender to greed, resorting to unethical means for financial gain. Generally, individuals with a greedy nature are often men.

(b) Example of different sentence structures with varied levels of context. Context gradually increases from S1 to S5

Figure 6: Word and Sentence examples for a study on Log Probability Bias method for Bangla

B Details of Benchmarking methods

We describe the methodological details of the benchmarking experiments we conduct for Bangla.

B.1 Word Embedding Association Test (WEAT)

In order to quantify bias in English language embeddings, Caliskan et al. (2017) proposed the Word Embedding Association Test (WEAT), they compared two sets of target words with two sets of attribute words. They quantified their comparison by calculating the effect size (d) and statistical significance (p -value).

To execute the WEAT experiment, we use distinct sets of **Target** vs **Attribute** words categorized as shown in Table 1. For the extraction of embedding vectors corresponding to each word, we use two distinct Bangla pretrained word embedding models, specifically the **word2vec** and the **GloVe** embedding models. Subsequently, we compute effect sizes and corresponding p -values to assess statistical significance, with a significance threshold set at $p < 0.05$.

To calculate the effect size -

$$ES = \frac{\text{mean}_{x \in X} s(x, A, B) - \text{mean}_{y \in Y} s(y, A, B)}{\text{std_dev}_{w \in X \cup Y} s(w, A, B)}$$

In this context, “effect size” refers to the size of the bias as determined by the WEAT metric. The effect size is expressed as Cohen’s d , a common measure of the difference between two means that has been standardized by the standard deviation of the data. A larger Cohen’s d value denotes a stronger bias between the targets. According to Cohen’s effect size metric, $d > |0.5|$ and $d > |0.8|$ are medium and large effect sizes, respectively. (Rice and Harris, 2005)

B.2 Sentence Encoder Association Test (SEAT)

To perform the SEAT experiment for Bangla language, following the approach of May et al. (2019), we curate a comprehensive list of sentence templates. Each target word from the WEAT target list is incorporated into the SEAT template sentences. We use the Bangla translated version of the semantically bleached templates in May et al. (2019). We use the final layer of BanglaBERT (Bhattacharjee et al., 2022) to extract embeddings for each sentence. We then use these embeddings to calculate the effect size of the curated list of sentences based on the mentioned categories.

C CEAT Data Extraction

Data extraction for CEAT experiment is carried out using the **Bangla2B+** dataset. Examples of naturally occurring sentences from the dataset is provided in Figure 5a. As mentioned in section 4.2, we utilize pattern-matching to extract the suitable sentences. We create unique suffix groups for each word from our categories.

C.1 Relevance of Suffix

Unlike English, Bangla words typically contain suffixes when used in a sentence. Due to this characteristic, it was necessary for our methodology to consider the presence of suffixes before applying pattern matching to extract the groups of sentences for each word in the category. Figure 7a depicts the presence of suffix for a root word in naturally occurring sentences. If the presence of suffixes is not accounted for, the number of extracted sentences for a specific root word is significantly reduced.

C.2 Significance of Unique Suffix Groups

Our investigation further reveals that relying on a common group of suffixes across all words in our **Target** vs **Attribute** categories is inadequate and error-prone. Since in Bangla, each word has its own unique set of suffixes that are appended in sentences. In Figure 7b, we provide examples of correct vs wrong sets of suffixes for a specific root word. It is evident that if wrong set of suffixes are applied to a word, it would result in erroneous extraction of sentences.

To tackle this characteristic, we create a list of **21** distinct suffix groups and link each word to its corresponding group. Each suffix group contains 2 to 15 suffixes based on the type of word the group will be assigned to. This process enables accurate sentence extraction based on the specific word-suffix combination within the dataset.

D CEAT

D.1 Random Effects Model

Random Effects Model⁵ (also known as Variance Components Model) is a statistical model where the model parameters are random variables. The model assumes that the data being analysed are drawn from a hierarchy of different populations whose differences relate to that hierarchy. Our

⁵“Random effects model”, Wikipedia, last modified 8 December, 2023, https://en.wikipedia.org/wiki/Random_effects_model

Language	Root Word	Suffix Group	Regex matching from Dataset (token length=3)
Bangla	মা	"কে", "য়েদের", "য়েরা", "য়ের"	আমি মাকে বলেছিলাম যে আজ আমার পরীক্ষার ফলাফল প্রকাশিত হবে।
			স্কুলে শিক্ষকরা বললেন যে মায়েরা সবসময় তাদের সন্তানদের খেয়াল রাখেন।
			শিশুদের ভালো মানসিক বিকাশের জন্য মায়েদের সঠিক দিকনির্দেশনা দেওয়া উচিত।
English	Mother	Not applicable	I told my mother that my exam results would be published today.
			The teachers at school said that mothers always take care of their children.
			For the good mental development of children, mothers should be given proper guidance.

(a) Presence of suffix in Bangla sentences for a specific root word.

Language	Root Word	Suffix Group	Regex matching from Dataset (token length=3)
Bangla	মা	Correct	"কে", "য়েদের", "য়েরা", "য়ের"
			আমি মাকে বলেছিলাম যে আজ আমার পরীক্ষার ফলাফল প্রকাশিত হবে।
		Wrong	"রা", "টি", "টির"
			স্কুলে শিক্ষকরা বললেন যে মায়েরা সবসময় তাদের সন্তানদের খেয়াল রাখেন।
English	Mother	Correct	Not Applicable
			I told my mother that my exam results would be published today.
		Wrong	Not Applicable
			The teachers at school said that mothers always take care of their children.

(b) Importance of unique suffix group for a specific root word.

Figure 7: Relevance and Uniqueness of suffix groups for CEAT Data Extraction

calculation of CEAT assumes that the differences between effect size due to contextualized variation for two sets of target words in terms of their relative similarity to two sets of attribute words is accounted for some random variable uncorrelated with independent variables.

Components of the same category have constant heterogeneity over time. These differences are caused by random contextual factors, represented by a random variable that does not directly influence the independent variables in the model.

The effect size (*Cohen's d*) for i^{th} sample is calculated by

$$ES_i = \frac{mean_{x \in X} s(x, A, B) - mean_{y \in Y} s(y, A, B)}{std_dev_{w \in X \cup Y} s(w, A, B)}$$

The in-sample variance estimation (denoted by V_i), is the square of $std_dev_{w \in X \cup Y} s(w, A, B)$. The between-sample variance, $\sigma_{between}^2$ is estimated using the same principle of ANOVA and calculated using the formula

$$\sigma_{between}^2 = \begin{cases} \frac{Q - (N-1)}{c} & \text{if } Q \geq N - 1 \\ 0 & \text{if } Q < N - 1 \end{cases}$$

where

$$W_i = \frac{1}{v_i},$$

$$c = \sum W_i - \frac{\sum W_i^2}{\sum W_i},$$

and

$$Q = \sum W_i ES_i^2 - \frac{\sum (W_i ES_i)^2}{\sum W_i}$$

Category	WEAT		SEAT	CEAT	Log Probability Bias
	(word2vec)	(GloVe)			
C1: Flowers/Insects (Pleasant/Unpleasant)	1.77*	1.27*	0.89*	1.225*	0.89*
C2: Music/Weapons (Pleasant/Unpleasant)	1.53*	0.99*	-0.03	-0.226*	0.42*
C3: Male/Female names (Pleasant/Unpleasant)	0.38	1.35*	0.78*	0.182*	0.22
C4: Male/Female names (Career/Family)	1.44*	-0.18	-0.58	0.639*	0.71*
C5: Male/Female terms (Career/Family)	0.42	0.17	-0.44	0.263*	0.62*
C6: Math/Art (Male/Female terms)	1.00*	0.68*	-0.17	0.258*	0.93*
C7: Math/Art (Male/Female names)	-0.17	-0.93	-0.67	-0.643*	0.48*
C8: Science/Art (Male/Female terms)	-0.22	-0.20	-0.76	0.366*	0.98*
C9: Science/Art (Male/Female names)	0.23	-1.03	-1.13	-0.591*	0.70*

Table 4: Effect size of bias measurements for various experiments (* indicates statistically significant at $p < 0.05$)

The weight v_i is the assigned weight to each effect size in measuring the *combined effect size* (CES). The parameter is determined by calculating the inverse of the sum of estimated in-sample variance V_i and estimated between-sample variance in the distribution of random effects $\sigma_{between}^2$.

$$v_i = \frac{1}{V_i + \sigma_{between}^2}$$

CES is the sum of the weighted effect sizes divided by the sum of all weights,

$$CES(X, Y, A, B) = \frac{\sum_{i=1}^N v_i ES_i}{\sum_{i=1}^N v_i}$$

Hypothesis Test: The standard error (SE) of CES is calculated to derive the hypothesis test. SE is calculated with the formula below

$$SE(CES) = \sqrt{\frac{1}{\sum_{i=1}^N v_i}}$$

Based on the central limit theorem, the limiting form of the distribution of $\frac{CES}{SE(CES)}$ is the standard normal distribution (Montgomery and Runger, 2010). We noticed that some of the CES values are negative, so we use a two tailed p - value which can test the bias significance in two directions. The hypothesis for which the two-tailed p - value provides significance test is that there is no difference between all the contextualized variations of the two sets of target words in terms of their relative similarity to two sets of attribute words is given by the following formula

$$P_{combined}(X, Y, A, B) = 2 \times [1 - \phi(|\frac{CES}{SE(CES)}|)]$$

where ϕ stands for the standard cumulative distribution function and SE stands for the standard error.

D.2 Language Models Used for CEAT Embedding Extraction

For extracting the embeddings necessary for CEAT experimentation, we use the output of the following models:

- **BanglaBERT Large** (Bhattacharjee et al., 2022) was trained using ELECTRA methodology (Clark et al., 2020). It contains 24 hidden layers. We use the outputs of the final layers as our word embeddings. We use both the *generator* (52M parameters) and the *discriminator* (339M parameters) versions of the model separately as they are trained on *Masked Language Modelling* (MLM) and *Replaced Token Detection* (RTD) objectives respectively.
- **MuRIL** (Khanuja et al., 2021) is a BERT (Devlin et al., 2019) model trained on two different language modelling objectives, *Masked Language Modelling* (MLM) and *Translation Language Modelling* (TLM). We use the MuRIL-large-cased version with 24 layers and 506M parameters. We extract the hidden unit values of the top layer as its CWE of 1024 dimensions. The base model has 238M parameters.
- **XLM-RoBERTa** (Conneau et al., 2020) is a transformer-based model (Vaswani et al., 2017) designed for multilingual natural language processing tasks. It was trained with multilingual MLM objective. We use the large version with 24 hidden layers and 560M parameters. The embeddings are taken from top-most layer with 1024 dimensions. The large model comprises of 560M parameters.

D.3 Results for sample size, $N = 5000$

In the main section, we mentioned a short result for CEAT analysis in Table 2. We present that in

broader form with more segment length variations in table 5. We include two more segment lengths for each model here.

D.4 Results for sample size, $N = 1000$

In this section, we are focusing on a smaller sample size, specifically $N = 1000$. We present our results in Table 6. One noticeable change is the reduction in cells demonstrating statistically significant values. However, for the most part, individual cells show only minor changes.

The key characteristics of the model, as highlighted for the $N = 5000$ sample in Table 5, are still quite evident in Table 6. This suggests that achieving similar results is possible even with a reduced sample size, especially when faced with resource constraints. Nevertheless, it is crucial to recognize that while the overall trends in the model's behavior remain consistent, there are nuanced alterations in the statistical significance of certain cells.

Given these observations, we propose that optimizing the sample size could be a promising avenue for further investigation. Determining the optimal sample size, one that ensures reliable results without sacrificing statistical significance, presents an interesting area for future research and in-depth exploration.

Category		BanglaBERT-L Generator				BanglaBERT-L Discriminator			
		9	25	75	>75	9	25	75	>75
C1: Flowers/Insects (Pleasant/Unpleasant)	f	1.225	1.208	1.206	1.207	0.551	0.466	0.46	0.461
	r	1.225	1.208	1.206	1.207	0.551	0.466	0.462	0.453
C2: Music/Weapons (Pleasant/Unpleasant)	f	-0.226	-0.239	-0.256	-0.258	-0.26	-0.18	-0.158	-0.168
	r	-0.226	-0.239	-0.256	-0.258	-0.271	-0.179	-0.16	-0.158
C3: Male/Female names (Pleasant/Unpleasant)	f	0.182	0.16	0.165	0.166	<0.001*	0.015*	0.017	0.018
	r	0.182	0.16	0.165	0.166	-0.003*	0.009*	0.023	0.018
C4: Male/Female names (Career/Family)	f	0.639	0.732	0.73	0.73	0.032	0.023	0.018	0.016*
	r	0.64	0.732	0.73	0.73	0.034	0.018	0.021	0.017
C5: Male/Female terms (Career/Family)	f	0.263	0.252	0.254	0.245	-0.036	-0.038	-0.03	-0.023*
	r	0.263	0.252	0.254	0.245	-0.041	-0.033	-0.029	-0.028
C6: Math/Art (Male/Female terms)	f	0.258	0.141	0.128	0.123	0.023*	0.061	0.069	0.063
	r	0.258	0.141	0.128	0.123	0.013*	0.066	0.066	0.063
C7: Math/Art (Male/Female names)	f	-0.643	-0.639	-0.638	-0.637	0.253	0.269	0.277	0.27
	r	-0.643	-0.639	-0.638	-0.637	0.255	0.266	0.264	0.288
C8: Science/Art (Male/Female terms)	f	0.366	0.287	0.269	0.273	-0.04	-0.036	-0.05	-0.046
	r	0.366	0.287	0.269	0.273	-0.052	-0.044	-0.056	-0.049
C9: Science/Art (Male/Female names)	f	-0.591	-0.643	-0.654	-0.651	-0.115	-0.145	-0.153	-0.142
	r	-0.591	-0.643	-0.654	-0.651	-0.125	-0.142	-0.159	-0.153

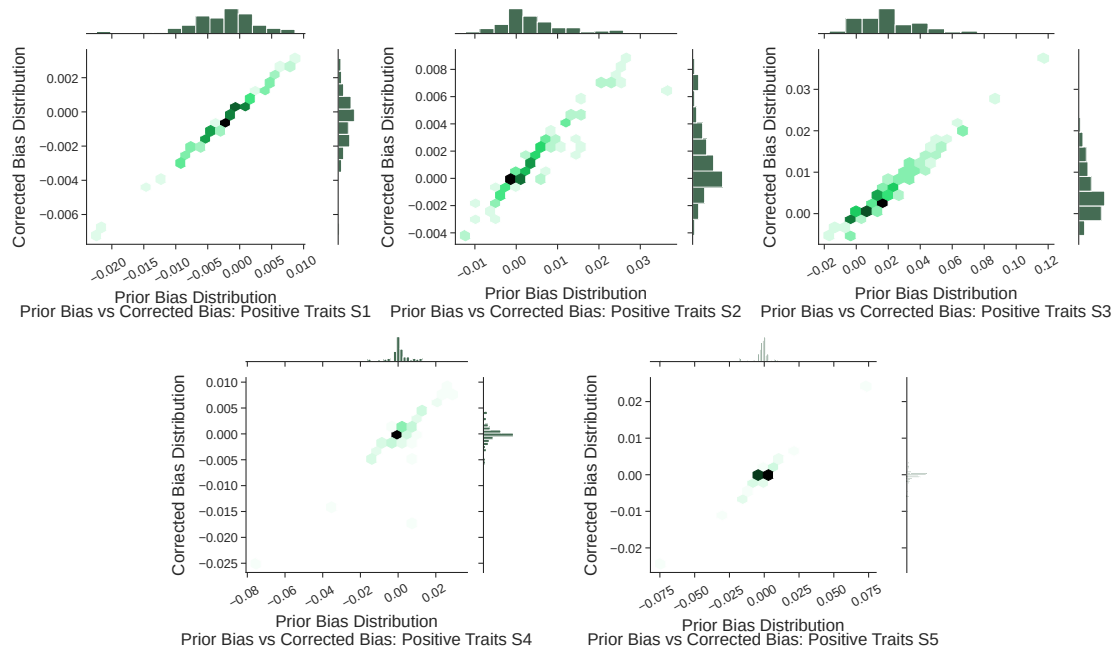
Category		MuRIL-Large(cased)				XLM-RoBERTa Large(cased)			
		9	25	75	>75	9	25	75	>75
C1: Flowers/Insects (Pleasant/Unpleasant)	f	1.193	1.215	1.2	1.208	0.279	0.478	0.493	0.495
	r	1.194	1.214	1.201	1.201	0.273	0.472	0.51	0.482
C2: Music/Weapons (Pleasant/Unpleasant)	f	0.477	0.372	0.379	0.377	0.566	0.738	0.769	0.767
	r	0.478	0.38	0.367	0.372	0.574	0.749	0.774	0.766
C3: Male/Female names (Pleasant/Unpleasant)	f	0.479	0.635	0.661	0.651	0.065	0.063	0.048	0.056
	r	0.484	0.63	0.649	0.661	0.064	0.05	0.059	0.04
C4: Male/Female names (Career/Family)	f	0.016	-0.085	-0.089	-0.085	-0.21	-0.277	-0.292	-0.286
	r	0.024	-0.092	-0.089	-0.083	-0.203	-0.263	-0.267	-0.326
C5: Male/Female terms (Career/Family)	f	0.226	0.256	0.267	0.264	-0.14	-0.199	-0.21	-0.21
	r	0.23	0.254	0.258	0.257	-0.152	-0.214	-0.21	-0.216
C6: Math/Art (Male/Female terms)	f	0.421	0.426	0.409	0.408	-0.103	-0.174	-0.143	-0.134
	r	0.414	0.432	0.421	0.417	-0.114	-0.174	-0.152	-0.133
C7: Math/Art (Male/Female names)	f	-0.16	-0.225	-0.272	-0.284	-1.26	-1.195	-1.191	-1.191
	r	-0.151	-0.23	-0.277	-0.27	-1.262	-1.202	-1.193	-1.194
C8: Science/Art (Male/Female terms)	f	-0.058	0.007*	0.008*	0.01	-0.076	-0.092	-0.104	-0.104
	r	-0.06	0.013*	0.006*	0.013	-0.08	-0.084	-0.075	-0.106
C9: Science/Art (Male/Female names)	f	-0.009*	-0.184	-0.213	-0.215	-0.296	-0.311	-0.313	-0.317
	r	-0.022	-0.183	-0.206	-0.204	-0.293	-0.307	-0.309	-0.312

Table 5: **Effect size of social bias measurements for different language models.** The bias is reported with overall magnitude of CES (d , with rounded values) and statistical significance with two-tailed p -values (p , significant at $p < 0.005$). The cells with *asterisk* (*) are statistically in-significant. The data results from CES pooling $N = 5000$ samples from random-effects model. The first row of each WEAT category uses fixed set of samples for each models, denoted as **f** and the second row uses completely random set of samples denoted as **r**. The light, medium and dark shades of grey are used to indicate small, medium and large effect size respectively.

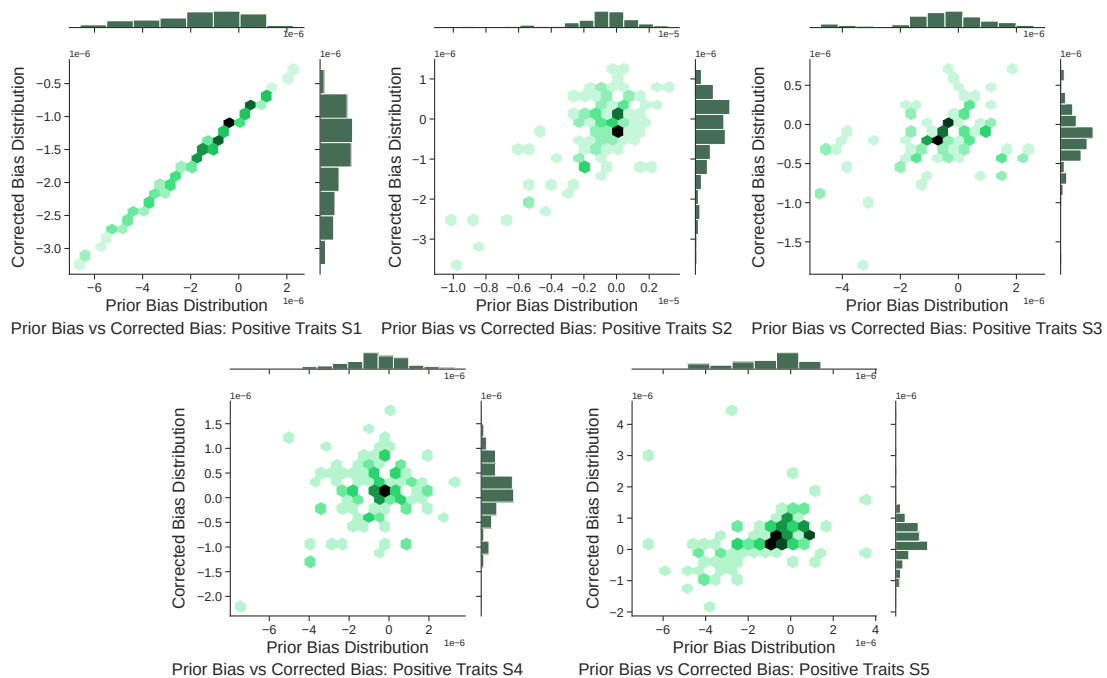
Category		BanglaBERT-Generator				BanglaBERT-Discriminator			
		9	25	75	>75	9	25	75	>75
C1: Flowers/Insects (Pleasant/Unpleasant)	f	1.226	1.206	1.212	1.22	0.532	0.419	0.461	0.5
	r	1.223	1.199	1.212	1.214	0.574	0.462	0.484	0.445
C2: Music/Weapons (Pleasant/Unpleasant)	f	-0.224	-0.249	-0.245	-0.256	-0.263	-0.159	-0.18	-0.147
	r	-0.232	-0.235	-0.254	-0.252	-0.246	-0.192	-0.174	-0.128
C3: Male/Female names (Pleasant/Unpleasant)	f	0.19	0.149	0.165	0.178	-0.018*	0.026*	0.022*	0.029
	r	0.18	0.15	0.166	0.167	-0.002*	0.005*	0.014*	0.028*
C4: Male/Female names (Career/Family)	f	0.643	0.734	0.723	0.726	0.019*	0.032	0.015*	0.007*
	r	0.641	0.751	0.735	0.745	0.016*	0.007*	0.012*	0.017*
C5: Male/Female terms (Career/Family)	f	0.259	0.245	0.248	0.243	-0.053*	-0.01*	-0.008*	-0.021*
	r	0.275	0.262	0.253	0.256	-0.056*	-0.048	0.001*	-0.058
C6: Math/Art (Male/Female terms)	f	0.254	0.176	0.15	0.14	0.023*	0.083	0.054	0.057
	r	0.262*	0.157	0.13	0.125	0.021	0.074	0.075	0.078
C7: Math/Art (Male/Female names)	f	-0.629	-0.642	-0.639	-0.628	0.252	0.294	0.273	0.295
	r	-0.621	-0.63	-0.63	-0.62	0.266	0.282	0.243	0.259
C8: Science/Art (Male/Female terms)	f	0.366	0.306	0.279	0.283	-0.041*	-0.042	-0.055*	-0.046
	r	0.375	0.277	0.282	0.305	-0.031	-0.051	-0.031	-0.082
C9: Science/Art (Male/Female names)	f	-0.593	-0.635	-0.647	-0.65	-0.124	-0.147	-0.167	-0.159
	r	-0.584	-0.635	-0.646	-0.638	-0.117	-0.143	-0.169	-0.163

Category		MuRIL-Large(cased)				XLM-RoBERTa Large(cased)			
		9	25	75	>75	9	25	75	>75
C1: Flowers/Insects (Pleasant/Unpleasant)	f	1.196	1.218	1.213	1.201	0.275	0.472	0.51	0.482
	r	1.192	1.207	1.195	1.201	0.278	0.469	0.499	0.494
C2: Music/Weapons (Pleasant/Unpleasant)	f	0.478	0.376	0.377	0.358	0.574	0.749	0.774	0.766
	r	0.469	0.375	0.382	0.363	0.572	0.733	0.769	0.769
C3: Male/Female names (Pleasant/Unpleasant)	f	0.48	0.626	0.653	0.648	0.064	0.05	0.059	0.04
	r	0.497	0.619	0.655	0.659	0.068	0.044	0.051	0.058
C4: Male/Female names (Career/Family)	f	0.015	-0.065	-0.08	-0.093	-0.203*	-0.263	-0.267	-0.326
	r	0.037	-0.082	-0.076	-0.1	-0.216	-0.279	-0.307	-0.302
C5: Male/Female terms (Career/Family)	f	0.205	0.274	0.273	0.245	-0.152	-0.214	-0.21	-0.216
	r	0.231	0.239	0.244	0.263	-0.158	-0.193	-0.202	-0.206
C6: Math/Art (Male/Female terms)	f	0.412	0.436	0.436	0.386	-0.114	-0.174	-0.152	-0.133
	r	0.411	0.442	0.414	0.419	-0.111	-0.224	-0.139	-0.106
C7: Math/Art (Male/Female names)	f	-0.141	-0.226	-0.284	-0.288	-1.262	-1.202	-1.193	-1.194
	r	-0.131	-0.233	-0.275	-0.276	-1.261	-1.2	-1.187	-1.187
C8: Science/Art (Male/Female terms)	f	-0.068	0.026	0.018	0.014	-0.08	-0.084*	-0.075*	-0.106*
	r	-0.055	0.006*	0.011*	0.011*	-0.076	-0.089	-0.093	-0.102
C9: Science/Art (Male/Female names)	f	-0.013*	-0.179	-0.231	-0.207	-0.293	-0.307	-0.309	-0.312
	r	-0.016*	-0.207	-0.21	-0.205	-0.291	-0.315	-0.309	-0.313


Table 6: **Effect size of social bias measurements for different language models.** The bias is reported with overall magnitude of CES (d , with rounded values) and statistical significance with two-tailed p -values (p , significant at $p < 0.005$). The cells with *asterisk* (*) are statistically in-significant. The data results from CES pooling $N = 1000$ samples from random-effects model. The first row of each WEAT category uses fixed set of samples for each models, denoted as **f** and the second row uses completely random set of samples denoted as **r**. The light, medium and dark shades of grey are used to indicate small, medium and large effect size respectively. Compared to pooling with $N = 5000$, more cells with statistically in-significant values are seen.



(a) Prior Bias Score vs Corrected Bias Score plots for positive traits in **BanglaBERT** - Large Generator.



(b) Prior Bias Score vs Corrected Bias Score plots for positive traits in **MuRIL** - Large (cased).



(c) Prior Bias Score vs Corrected Bias Score plots for negative traits in **MuRIL** - Large (cased).

Figure 8: A comparison between model behaviors for different sentence structures in Log Probability Bias Score Test.