Evaluation of Gender Bias in Amharic Word Embedding Model

Beimnet Zenebe[1][0009-0004-3952-7851], Solomon Gizaw[1], and Yalemisew Abgaz[2][0000-0002-3887-5342]

[1] Addis Ababa University, Addis Ababa, Ethiopia
[2] Dublin City University, Dublin, Ireland
Beimnet.Girma,Solomong@aau.edu.et, YalemisewM.Abgaz@dcu.ie

Abstract. Bias in natural language processing systems can perpetuate and exacerbate societal inequalities, reflecting and potentially amplifying existing biases in human language and culture. Amharic, as the official language of Ethiopia, holds cultural and linguistic significance, making it imperative to assess potential biases within its computational representations. This research paper investigates the presence and extent of gender bias in Amharic text corpora. The research utilizes gendered word pairs to capture gender representation in the word embeddings and quantifies the degrees of gender bias present in profession words. We found that profession words carried stereotypical implicit biases with most occupations leaning towards male. Profession words like "nurse" and "house-maid" align with societal gender dynamics, displaying significant female associations. Additionally, professions in the arts and athleticism demonstrate a robust female-leaning bias, while physically demanding and educated professional roles tend to exhibit male-leaning biases. The study contributes insights into the gender dynamics encoded within the Amharic language informing strategies to reduce bias and fostering fair and unbiased representations for improved societal and technological outcomes.

Keywords: Gender Bias · Natural Language Processing · Word Embeddings · Bias Measurement · Amharic Corpus

1 Introduction

Natural Language Processing (NLP) is the heart of many Artificial Intelligence (AI) applications. NLP is vital in many areas, such as text translation, speech recognition, caption generation, and sentiment analysis. It is also a significant component of automated decision-making or support systems such as job application filtering, load authorization, and criminal verdict recommendations. However, despite the many applications of NLP and its vast, arguably positive impact on our lives, it is plagued with several issues; one of the most prominent ones being the presence of different societal (racial, gender, religious) biases [3,6,17]. Manzini et al. [12] found that Machine Learning (ML) models trained on Reddit data formed the relation "Black is to criminal as Caucasian is to police". Another study [23] also found that AI models for processing hate speech were 1.5 more likely to flag tweets as "hate speech" when written by African Americans or used African American Vernacular English (AAVE).

Gender bias is an unfounded preference or association of one gender over another. Gender bias is exhibited in NLP in different forms. A paper by Bolukbasi et al. titled "Man is to Programmer, Woman is to Homemaker" showed the problematic titular relation on pre-trained models [5]. Another study on machine translation showed that translating the English statement "She is a doctor. He is a nurse" into Hungarian (a gender-neutral language) and back into English again resulted in the statement "He is a doctor. She is a nurse" [9]. Many more examples [11,18,24] highlight the extent to which gender bias is present in NLP. Using these biased systems in downstream applications could result in biased decisions that may have real-world consequences.

The effects of gender bias in NLP have been well-researched for English and other resource-rich languages like French and Spanish. However, studies on Amharic, a low-resourced language and the most widely spoken in Ethiopia, are scarce. Despite its prevalence and status as the federal government's working language, Amharic lacks publicly available text data, pre-trained models, and NLP tools [2]. Its unique alphabet and complex morphology add to the computational challenges. Words can be generated from their base form by adding prefixes, suffixes, infixes, and circumfixes. Amharic nouns, adjectives, and verbs can be affixed to represent gender, number, tense, ownership, and many other dynamics [27]. The Amharic alphabet also contains various characters that make the same sound which are effectively redundant in modern writing and do not lend themselves to semantic differences [26]. Recently, efforts have been made to address these gaps, producing scientific corpora and other NLP resources.

This paper presents a pre-trained word embedding model for the Amharic language and tests the gender bias present in the model using existing approaches on other high and low-resource languages [13]. This study investigates the presence of gender bias in Amharic word embeddings by applying existing gender bias detection approaches. Specifically, the paper investigates, 1) the representation of gender in Amharic language, 2) the implicit association between occupation words and gender, and 3) the magnitude of bias across professions.

The rest of the paper is organized as follows: Section 2 describes previous studies on gender bias that guided this research. Section 3 delves into how we adopted and modified the existing methodology for Amharic copra. In Section 4, we present our findings on the gender distribution and bias measured in the model. Finally, we conclude this paper by discussing the results in Section 5.

2 Related Works

In their highly influential paper, Bolukbasi et al. [5] proposed a framework to measure gender bias in word embedding models such as Word2Vec [16], GloVe [19], and fastText [4]. Word embeddings, a significant breakthrough in deep learning [15], are high-dimentional vector representations of text that capture semantic meaning and relationships between words. Words with similar meanings will have vectors that are closer to each other. And more importantly for our research, these resulting vectors mathematically represent relationships between words. For example, the vector operation "king - man + woman" results in the vector for the word "queen." This capability aids various NLP tasks but can also amplify biases. As demonstrated in [5] research, taking the vector representation of a computer programmer, subtracting the vector for man (maleness), and adding the vector for woman (femaleness) resulted in the vector for homemaker, which is a concerning finding.

Bolukbasi et al. [5] introduced a framework to measure gender bias in word embedding models like Word2Vec, GloVe, and fastText. Word embeddings, significant in deep learning, are vector representations of text that capture semantic meaning and relationships between words. For example, the vector operation “king - man + woman” results in “queen.” This capability aids various NLP tasks but can also amplify biases. Bolukbasi’s study revealed that altering the “computer programmer” vector by gender led to "homemaker," highlighting concerning biases in embeddings. The authors defined gender bias as the unequal distance between vectors of gender-neutral words such as nurse, doctor, and programmer and vectors for gender-specific terms such as man, woman, son, and daughter. When a word devoid of gender implications semantically is closer in distance to one gender over the other, this represents bias. Instead of using individual words to represent the gender spectrum, the paper uses several pairs of gender-specific words (e.g., she-he, her-him, queen-king) to identify a gender direction in which other words could be measured.

Matthews et al. [13] applied the gender bias measurement method from [5] to Chinese, Spanish, Arabic, German, French, Farsi, Urdu, and Wolof using Wikipedia corpora. They identified seven gender-specific word pairs common to all nine languages and used principal component analysis (PCA)to find a gender direction. Starting with 337 profession words identified in [5], they narrowed it to 32 diverse professions. Their methodology worked across languages, though for some, like Chinese and Arabic, the gender direction was less pronounced, suggesting the word sets might not fully capture gender. Sakai and Suzuki [22] used Bolukbasi’s method to study gender bias in Japanese word-embedding models. They identified gendered antonym pairs, calculated the gender subspace, and analyzed top occupations from the Japanese National Census. They found female bias in traditionally female-dominated occupations like nurse, midwife, and hairdresser, with these words being more biased than their male counterparts. The gender bias in embeddings reflected workforce gender percentages.

Caliskan et al. [7] analyzed gender biases in word embedding models using the Single-Category Word Embedding Association Test. They analyzed associated gender bias in word frequency, part-of-speech tags, semantic categories, and valence, arousal, and dominance. They found that 77% of the 1,000 most frequent words were more associated with men, indicating a masculine default. They also discovered male-associated words were typically action-oriented verbs, while female-associated words were often found to be adjectives and adverbs related to emotion and appearance. Similarly male-associated words pertained to tech, engineering, religion, sports, and violence, whereas female-associated concepts included slurs, sexual content, and domestic terms.

Wairagala et al. [25] showed that word embeddings introduced gender bias in gender-neutral languages like Luganda. They built a Luganda-English translator and found that gender-specific pronouns were introduced in English translations, with female pronouns linked to family and male pronouns to careers.

To the best of our knowledge, no comprehensive study has been conducted on the presence and magnitude of gender bias in Amharic language processing. Nevertheless, studies have been conducted on the manifestation of linguistic sexism in Ethiopian local languages. Raga [20] examined sexism and gender bias in Amharic, Affan Oromo, and Gamo. Using data from native speakers, the study found deep-rooted sexism in these languages, reflecting and reinforcing male dominance in their societies. The research elucidated how words for heroism and courage were linked to the masculine gender, while those for the female gender conveyed cowardice and helplessness.

3 Methodology

3.1 Data Collection

The initial phase of this study involved acquiring Amharic texts to train a word embedding model. Due to limited prior literature, empirical assessments of word embedding models in the context of the Amharic language were necessary. We used the Contemporary Amharic Corpus (CACO) [10], an open-source dataset of 25,199 documents from newspapers, magazines, fiction, historical novels, political books, children's books, and the Amharic Bible. After preprocessing for standardization, the corpus contained 24 million tokens, 1,605,451 sentences, and totaled 280MB.

3.2 Preprocessing

After obtaining the COCA corpus, we refined the dataset, which had already undergone spelling correction, punctuation normalization, and sentence extraction [10]. We normalized common orthographic variations in words used for gender-bias analysis. In Amharic, the same word can have multiple spellings due to its phonetic nature, such as"her" written as "እሷ" or "እርሷ." Homophonic characters also allow spelling variations of the same word like "him" written as "እሱ," "ዕሱ," "ዕርሱ," or "እርሱ", all read the same way. We conducted word count analyses on these variations and transformed them into their canonical spellings to consolidate and normalize their representations.

3.3 Word Embedding

We used the Python Gensim library [21] to train a word embedding model based on the corpus. We identified sentences in the corpus using the Amharic sentence end marker “ ። ” and used the space boundary to identify token words. We applied the CBOW algorithm using Word2Vec model running a vector size of 300 and 30 epoch with minimum count of 10. We the tested the functionality of the model by inspecting the vector representation of similar words.

3.4 Gender Identification

Once a working word embedding model had been trained, we identified metrics for measuring gender bias. Bolukbasi et al. [5] proposed a method to assess gender bias by classifying words as gender-neutral (e.g., “umbrella,” “taxi,” “pilot”) or gender-specific (e.g., “king,” “queen,” “mother,” “father”). They used pairs of gender-specific words with similar meanings but differing in gender, like [mother-father] and [king-queen], to construct a gendered vector space, highlighting distinctions within these pairs.

Guided by [5,13], we set out to identify a set of gender-specific pairs that would be able to represent gender in the Amharic language. [5] identified ten English gender-specific pairs. Constrained by the fact that Bolukbasi et al.'s pairs did not have a direct translation in some of the languages that they were working on, [13] proposed a new set of gender-specific words to be used in all nine languages. We investigated both gender-specific word sets and removed ones that either did not have meaningful translations in Amharic or were culturally irrelevant. This resulted in a set containing nine pairs of gender-specific words (her-him, woman-man, herself-himself, daughter-son, mother-father, queen-king, wife-husband, Mrs-Mr, hers-his). We have presented the gender-specific Amharic pairs and their possible orthographic representations along with their English translations in table 1.

Table 1. Identified gender-pair words

Amharic Pair	English Pair	Amharic Pair	English Pair
እሷ፣ እሱ	Her, Him	ንግሥት፣ ንጉሥ	Queen, King
ሴት፣ ወንድ	Woman, Man	ሚስት፣ ባል	Wife, Husband
ራሷ፣ ራሱ	Herself, Himself	ወይዘሮ፣ አቶ	Mrs, Mr
ሴት ልጅ፣ ወንድ ልጅ	Daughter, Son	የሷ፣ የሱ	Hers, His
እናት፣ አባት	Mother, Father		

3.5 Identifying Gender Direction

As mentioned earlier, bias is measured by the distance between a gender-neutral word and gender-specific words of similar meaning but different gender. However, factors like polysemy and word count randomness make it more robust to use an aggregate of identified gender words to measure gender bias. Bolukbasi et al. [5] suggest using PCA on the vectors of gender-specific word pairs to identify the gender direction. The first eigenvalue is expected to capture the main aspect of gender in the corpus due to the highly gendered nature of the chosen word pairs.

We performed PCA on the gender-specific word set defined above and picked the explained variance ratio of the first component, 0.24, to be our gender direction. The first component value was not more significant than the following values (0.13 and 0.11). We experimented with changing the gendered word set to maximize this value, to no success. We theorized that since the Amharic language relies heavily on morphing root words to denote gender, just using the gendered pairs might not be enough to define and capture gender. For this experiment, we went ahead with the first value of the PCA as the gender direction.

3.6 Measuring Gender Bias

Bolukbasi et al. [5] proposed a formula to measure a specific word's bias (location) on the gender subspace. Given a set of gender-neutral occupation words (N) and a gender direction (g), they geometrically defined direct gender bias in a training corpus as:

$$\frac{1}{|N|}\sum_{w\epsilon N}|cos(w,g)|^{c} \tag{1}$$

Where c is how strictly we want to measure bias. Setting c to 0, we essentially will only accept a gender bias of 0 from the formula if and only if w has no overlap with g. This might be desirable in settings sensitive to gender considerations, such as job hirings or student admissions. Setting c to 1 allows for a more lenient bias measurement. This formula can be used to determine the gender bias score of individual words.

3.7 Confirming Gender Direction

To reaffirm that the gender direction accurately captures gender in the corpus, we measured the gender "bias" of the gender-specific sets. The "bias" in this case is whether or not we could accurately identify words in this set as belonging to either male or female. Using *Equation 1*, we calculated the bias score of each word (w) in the gender-specific set (N) with our gender direction set as 0.24 (g). Figure 1 shows the gender bias of the starting gender-specific word sets. All male-gendered words were found to be male-leaning, and female-gendered words were found to be female-leaning. We considered this as a positive sign that the identified gender direction was functionally adequate.

3.8 Defining Profession Words Set

To identify gender-neutral profession words, we started with Matthew et al.'s set of 32 professions, removing those absent from the corpus (e.g., adventurer, filmmaker, astronaut). We added common occupations with significant corpus presence and included multiple Amharic translations for some professions. For

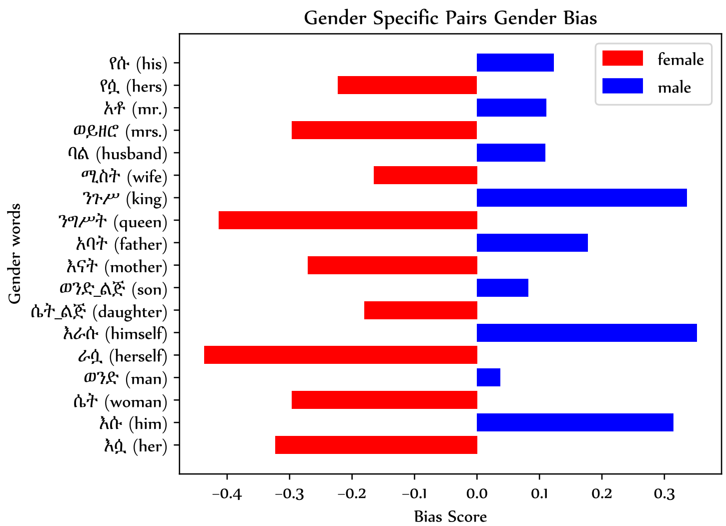

Fig. 1. Gender bias scores of the identified gender-specific word pairs

example, we included variants of “teacher,” “engineer,” “scientist,“driver,” and “farmer,” as well as both translated and phonetically spelled forms of “doctor” (e.g., ዶክተር and ሃኪም). This resulted in a list of 38 occupation words. Table 2 shows these professions and their corpus frequencies.

4 Results

After having identified the profession words we wanted to evaluate, we used Equation 1 to measure the gender inclination of each word. Figure 2 shows the gender bias of the identified occupation words.

We found that most professions were male-leaning. This could be due to the fact in Ethiopia, men dominate the majority of the workforce; therefore most professions are biased in favor of the male gender. There were some exceptions; for instance, the word for house-help or maid is highly biased towards female. This reflects the reality of the culture where the majority of house helpers are women. Nurse was heavily female leaning as expected. A few professions were female-leaning, but their gender bias score towards female was minimal.

5 Discussion

Our analysis shows that Amharic word embeddings reflect stereotypical gender norms, with 76.3% (29/39) of selected occupation words biased toward males.

Table 2. Identified profession words and their frequency in the corpus

Profession	**Word Count**	**Profession**	**Word Count**
Nurse (ነርስ)	124	Soldier (ወታደር)	796
Teacher (አስተማሪ፣ መምህር)	739 + 1584	Journalist (ጋዜጠኛ)	1448
Writer (ደራሲ፣ ጸሐፊ)	1152 + 596	Student (ተማሪ)	2279
Engineer (መሐንዲስ፣ ኢንጂነር)	98 + 564	Athlete (አትሌት)	2522
Scientist (ሳይንቲስት፣ ተመራማሪ)	149 + 403	Actor (ተዋናይ)	445 + 21
Manager (ሥራ አስኪያጅ፣ ዳይሬክተር)	974 + 34 + 939	Governor (ገዥ)	566 + 52
Driver (ሾፌር፣ አሽከርካሪ)	106 + 210 + 240	Farmer (ገበሬ፣ አርሶ አደር)	579 + 1102
Person (ሰዉ)	161	Musician (ሙዚቀኛ፣ ዘፋኝ)	157 + 139
Lawyer (ጥበቃ)	478	Artist (አርቲስት)	879
Aide (ረዳት)	795	Ambassador (እንደራሴ)	88
Judge (ዳኛ)	950	Analyst (ተንታኝ)	90
Advisor (አማካሪ)	722	Comedian (ኮሜዲያን)	89
Maid (የቤት ሠራተኛ)	49	Worker (ሠራተኛ)	1332
Priest (ቄስ)	320	Doctor (ሐኪም፣ ዶክተር)	625 +3904
Guard (ዘበኛ)	186	Janitor (ጽዳት)	131

The term ”nurse” had the strongest female association, consistent with other languages. Female-leaning embeddings were also found for roles like ”house-help” and ”janitors,” typically low-paying jobs requiring minimal education. This bias mirrors societal expectations and highlights disparities in employment opportunities and working conditions, underscoring the nuanced impact of language representations in Amharic.

Our study found a gender bias in embedding for the word “doctor” written in Amharic letters, unexpectedly leaning toward female associations, while the Amharic term ሃኪም showed a male-leaning bias. The reason for this discrepancy is unclear. We also observed that words related to the arts, such as “artist”, “singer”, “actor”, and “athlete”, showed a strong female-leaning bias. In contrast, many profession-related words, especially those for physically demanding jobs like “guard”, “farmer”, and “soldier”, and those requiring higher education like “lawyer”, “ambassador”, and “teacher” had male-leaning biases. This reflects societal dynamics where the most educated workforce is predominantly male. Additionally, we analyzed the terms “person” and “student”, which, though not professions, showed a male-leaning bias. It is crucial to note that in Amharic, gender-specific suffixes are used for most words, but the male form is often seen as neutral or inclusive. This contributes to the observed gender biases.

6 Conclusion

In this paper, we extended a famous methodology proposed by Bolukbasi et al [5] and experimented on by other researchers to measure gender bias in the Amharic language. We used Amharic gender words to capture gender in the corpus and quantified how certain occupation words related to each gender.

Our findings show that gender stereotypes and inequality in the workforce are reflected in word embeddings for occupation words, affecting the fairness of applications where gender should be irrelevant. The nuances of gender represen-

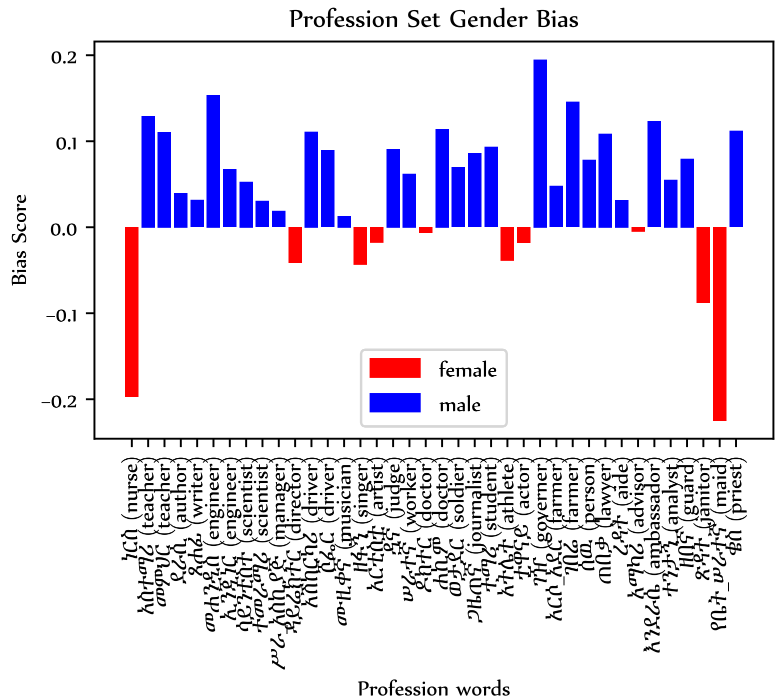

Fig. 2. Gender bias scores of selected professions

tation in the Amharic language, which extends beyond the binary distinction of male and female word pairs, pose an additional layer of complexity. Future research could delve into these unique gender markers to better understand their impact on natural language processing.

Beyond inherent biases in word embeddings, further study could involve surveying societal perceptions of gender associations with specific occupations. Developing Amharic lexical resources with gender annotations using semantic web techniques, like the Ontolex model [1,8,14], could enhance the overall understanding of cultural biases and strengthen our analysis.

Our research highlights the importance of addressing and mitigating gender bias in language representations to ensure fair and unbiased outcomes in various applications. As technology continues to shape societal norms, an ongoing commitment to refining these models becomes imperative, emphasizing the need for interdisciplinary collaboration between linguists, ethicists, and technologists.

References

1. Abgaz, Y.: Using OntoLex-lemon for representing and interlinking lexicographic collections of Bavarian dialects. In: Ionov, M., McCrae, J.P., Chiarcos, C., De-

clerck, T., Bosque-Gil, J., Gracia, J. (eds.) Proceedings of the 7th Workshop on Linked Data in Linguistics (LDL-2020). pp. 61–69. European Language Resources Association, Marseille, France (May 2020)

2. Argaw, A.A., Asker, L.: An Amharic stemmer: reducing words to their citation forms. In: Proceedings of the 2007 Workshop on Computational Approaches to Semitic Languages Common Issues and Resources - Semitic '07. p. 104. Association for Computational Linguistics, Prague, Czech Republic (2007). https://doi.org/10.3115/1654576.1654594
3. Birhane, A.: Algorithmic injustice: a relational ethics approach. Patterns **2**(2), 100205 (2021). https://doi.org/https://doi.org/10.1016/j.patter.2021.100205
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. Transactions of the Association for Computational Linguistics **5**, 135–146 (2016)
5. Bolukbasi, T., Chang, K.W., Zou, J.Y., Saligrama, V., Kalai, A.T.: Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In: Advances in Neural Information Processing Systems. vol. 29. Curran Associates, Inc. (2016)
6. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: Friedler, S.A., Wilson, C. (eds.) Proceedings of the 1st Conference on Fairness, Accountability and Transparency. Proceedings of Machine Learning Research, vol. 81, pp. 77–91. PMLR (23–24 Feb 2018)
7. Caliskan, A., Ajay, P.P., Charlesworth, T., Wolfe, R., Banaji, M.R.: Gender Bias in Word Embeddings: A Comprehensive Analysis of Frequency, Syntax, and Semantics. In: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society. pp. 156–170 (Jul 2022). https://doi.org/10.1145/3514094.3534162, arXiv:2206.03390 [cs]
8. Dorn, A., Wandl-Vogt, E., Abgaz, Y.M., Santos, A.B., Therón, R.: Unlocking cultural conceptualisation in indigenous language resources:collaborative computing methodologies. In: International Conference on Language Resources and Evaluation (2018)
9. Douglas, L.: AI is not just learning our biases; it is amplifying them. (Dec 2017)
10. Gezmu, A.M., Seyoum, B.E., Gasser, M., Nürnberger, A.: Contemporary Amharic Corpus: Automatically Morpho-Syntactically Tagged Amharic Corpus. In: Machonis, P., Barreiro, A., Kocijan, K., Silberztein, M. (eds.) Proceedings of the First Workshop on Linguistic Resources for Natural Language Processing. pp. 65–70. Association for Computational Linguistics, Santa Fe, New Mexico, USA (Aug 2018)
11. Hendricks, L.A., Burns, K., Saenko, K., Darrell, T., Rohrbach, A.: Women also snowboard: Overcoming bias in captioning models. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) Computer Vision – ECCV 2018, pp. 793–811. Springer International Publishing, Cham (2018)
12. Manzini, T., Yao Chong, L., Black, A.W., Tsvetkov, Y.: Black is to Criminal as Caucasian is to Police: Detecting and Removing Multiclass Bias in Word Embeddings. In: Burstein, J., Doran, C., Solorio, T. (eds.) 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). pp. 615–621. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). https://doi.org/10.18653/v1/N19-1062
13. Matthews, A., Grasso, I., Mahoney, C., Chen, Y., Wali, E., Middleton, T., Njie, M., Matthews, J.: Gender Bias in Natural Language Processing Across Human Languages. In: Pruksachatkun, Y., Ramakrishna, A., Chang, K.W., Krishna, S.,

Dhamala, J., Guha, T., Ren, X. (eds.) Proceedings of the First Workshop on Trustworthy Natural Language Processing. pp. 45–54. Association for Computational Linguistics, Online (Jun 2021). https://doi.org/10.18653/v1/2021.trustnlp-1.6
14. McCrae, J.P., Gil, J.B., Gràcia, J., Bitelaar, P., Cimiano, P.: The ontolex-lemon model: Development and applications (2017)
15. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality (2013)
16. Mikolov, T., Yih, W.t., Zweig, G.: Linguistic Regularities in Continuous Space Word Representations. In: Vanderwende, L., Daumé III, H., Kirchhoff, K. (eds.) Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 746–751. Association for Computational Linguistics, Atlanta, Georgia (Jun 2013)
17. Nair, S.: How biases in language get perpetuated by technology (Jan 2019)
18. Park, J.H., Shin, J., Fung, P.: Reducing Gender Bias in Abusive Language Detection. In: Riloff, E., Chiang, D., Hockenmaier, J., Tsujii, J. (eds.) Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. pp. 2799–2804. Association for Computational Linguistics, Brussels, Belgium (Oct 2018). https://doi.org/10.18653/v1/D18-1302
19. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Conference on Empirical Methods in Natural Language Processing (2014)
20. Raga, A.: Comparative analyses of linguistic sexism in Afan Oromo, Amharic, and Gamo. International Journal of Sociology and Anthropology **7**(9), 189–196 (Sep 2015). https://doi.org/10.5897/IJSA2014.0590, publisher: Academic Journals
21. Rehurek, R., Sojka, P.: Gensim – statistical semantics in python (2011)
22. Sakai, S., Suzuki, Y.: Evaluation of Gender Bias in Japanese Word Embedding Model (2023)
23. Sap, M., Card, D., Gabriel, S., Choi, Y., Smith, N.A.: The Risk of Racial Bias in Hate Speech Detection. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 1668–1678. Association for Computational Linguistics, Florence, Italy (2019). https://doi.org/10.18653/v1/P19-1163
24. Tatman, R.: Gender and Dialect Bias in YouTube's Automatic Captions. In: Hovy, D., Spruit, S., Mitchell, M., Bender, E.M., Strube, M., Wallach, H. (eds.) Proceedings of the First ACL Workshop on Ethics in Natural Language Processing. pp. 53–59. Association for Computational Linguistics, Valencia, Spain (Apr 2017). https://doi.org/10.18653/v1/W17-1606
25. Wairagala, E.P., Mukiibi, J., Tusubira, J.F., Babirye, C., Nakatumba-Nabende, J., Katumba, A., Ssenkungu, I.: Gender bias Evaluation in Luganda-English Machine Translation. In: Duh, K., Guzmán, F. (eds.) Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track). pp. 274–286. Association for Machine Translation in the Americas, Orlando, USA (Sep 2022)
26. Woldeyohannis, M., Meshesha, M.: Usable Amharic Text Corpus for Natural Language Processing Applications. Applied Corpus Linguistics **2**, 100033 (Aug 2022). https://doi.org/10.1016/j.acorp.2022.100033
27. Yeshambel, T., Mothe, J., Assabie, Y.: Learned Text Representation for Amharic Information Retrieval and Natural Language Processing. Information **14**(3), 195 (Mar 2023). https://doi.org/10.3390/info14030195, number: 3 Publisher: Multidisciplinary Digital Publishing Institute