

# Mitigating Gender Stereotypes in Hindi and Marathi

Neeraja Kirtane\*

Manipal Institute of Technology  
Manipal, India  
kirtane.neeraja@gmail.com

Tanvi Anand

University of Texas, Austin  
Texas, USA  
tanvianand@gmail.com

## Abstract

As the use of natural language processing increases in our day-to-day life, the need to address gender bias inherent in these systems also amplifies. This is because the inherent bias interferes with the semantic structure of the output of these systems while performing tasks in natural language processing. While research is being done in English to quantify and mitigate bias, debiasing methods in Indic Languages are either relatively nascent or absent for some Indic languages altogether. Most Indic languages are gendered, i.e., each noun is assigned a gender according to each language’s rules of grammar. As a consequence, evaluation differs from what is done in English. This paper evaluates the gender stereotypes in Hindi and Marathi languages. The methodologies will differ from the ones in the English language because there are masculine and feminine counterparts in the case of some words. We create a dataset of neutral and gendered occupation words, emotion words and measure bias with the help of Embedding Coherence Test (ECT) and Relative Norm Distance (RND). We also attempt to mitigate this bias from the embeddings. Experiments show that our proposed debiasing techniques reduce gender bias in these languages.

## 1 Introduction

Word embeddings are used in most natural language processing tools. Apart from capturing semantic information, word embeddings are also known to capture bias in society (Bolukbasi et al., 2016). While most research has been focused on languages like English, less research has been done on low-resource languages and languages that have a grammatical gender (Zhou et al., 2019). A language with grammatical has a gender associated with every noun irrespective of whether the noun is animate or inanimate, e.g., a river in Hindi has

feminine gender. In contrast, words like writer have masculine and feminine counterparts. This gender association affects the pronouns, adjectives, and verb forms used during sentence construction. Grammatical genders in Hindi are masculine and feminine. In Marathi, there additionally exists a third neutral gender as well. Spoken by more than 600 million people, Hindi is the 3rd most spoken language in the world. Marathi is the 14th most spoken language with approximately 120 million speakers<sup>1</sup>. Given the expanse and the amount of people speaking these languages, it is essential to address the bias introduced by the computational applications of these languages.

We create a dataset of occupations and emotions. The occupation dataset consists of gendered and neutral occupation titles. The emotion dataset has words of different emotions like anger, fear, joy, and sadness. First, we identify existing gender bias by defining a subspace that captures the gender information. There are several ways to find this information. We use Principal Component Analysis (PCA) and Relational Inner Product Association (RIPA) (Ethayarajh et al., 2019). We use the existing metrics for evaluation: Embedding Coherence Test (Dev and Phillips, 2019), Relative Norm Distance (Garg et al., 2018). We modify these formulas so that they are correctly applicable to these gendered languages. We perform our experiments on the FastText word embeddings.

Next, we mitigate the gender bias found by the aforementioned using two approaches: Projection and Partial Projection. In summary, the key contributions of this paper are:

1. Dataset of emotions, and gendered and neutral occupations in Hindi and Marathi.
2. Methods to quantify the bias present in Hindi and Marathi word embeddings using the

---

<sup>1</sup><https://www.mentalfloss.com/article/647427/most-spoken-languages-world>

\*First author

above dataset.

3. Mitigate the bias through existing debiasing techniques.

## 2 Related work

Previous work to quantify and measure bias was done by Bolukbasi et al. (2016). They tried to find out a gender subspace by using gender-definition pairs. They proposed a hard de-biasing method that identifies the gender subspace and tries to remove its components from the embeddings.

The majority amount of research on gender bias is being done in English, which is not gendered (Stanczak and Augenstein, 2021). Languages like Spanish or Hindi have a grammatical gender, i.e., every noun is assigned a gender. Zhou et al. (2019) was one of the first papers to examine bias in languages with grammatical gender like French and Spanish. They used a modified version of the Word Embedding Association Test (WEAT) (Caliskan et al., 2017) to quantify the bias.

Sun et al. (2019) suggested mitigation techniques to remove gender bias like data augmentation, gender-swapping, and hard de-biasing according to the downstream task in NLP.

Being low-resource languages, there is less research done in languages like Hindi and Marathi. Previous work in Indic Languages was done by Pujari et al. (2019) where they built an SVM classifier to identify the bias and classify it. The problem with this method is that it needs a labeled gender dataset beforehand to train the classifier. Recent work by Ramesh et al. (2021) tries to find out bias in English-Hindi machine translation. They implement a modified version of the TGBI metric based on grammatical considerations for Hindi. TGBI metric is used to detect and evaluate gender bias in Machine Translation systems. Malik et al. (2021) measure Hindi specific societal biases like religion bias and caste bias along with gender bias.

## 3 Data

In Bolukbasi et al. (2016), the authors have compiled a list of professions in English and tried to find bias in them. Similarly, we compile a list of 166 professions, each in Hindi and Marathi languages. We split the professions into two parts, first is gender-neutral  $P_{neu}$  and the other is gendered  $P_{gen}$ . Similarly, we create a list of words of different emotions similar to the one in Kiritchenko

and Mohammad (2018) in Hindi and Marathi languages. The emotions are broadly classified into four types: anger, fear, joy, and sadness.

We have verified this data with the help of 2 independent native speakers of these languages. We also create pair of feminine and masculine seed word pairs in both the languages to identify the gender subspace. For example: queen, king. We call them target words  $T$ . Target words for Hindi Language is shown in figure 1. The dataset is available here <sup>2</sup>

["महिला", "पुरुष"] ["woman", "man"]  
["बेटी", "बेटा"] ["daughter", "son"]  
["मां", "पिता"] ["mother", "father"]  
["लड़की", "लड़का"] ["girl", "boy"]  
["रानी", "राजा"] ["queen", "king"]  
["पत्नी", "पति"] ["wife", "husband"]

Figure 1: Example of seed words in Hindi and their English translation

We test the bias on our data using FastText embeddings. FastText is a word embedding method that extends the word2vec model. Instead of learning vectors for words directly, FastText represents each word as an n-gram of characters. This helps capture the meaning of shorter words and allows the embeddings to understand suffixes and prefixes. A skip-gram model is trained to learn the embeddings once the word has been represented using character n-grams (Bojanowski et al., 2016).

Morphology is the field of linguistics that studies the internal structure of words. Morphologically rich languages refer to languages that contain a substantial amount of grammatical information (Comrie, 1999). Indic languages are morphologically rich because of the existence of a large number of different word forms. FastText embeddings are the best choice for Indian Languages as they are capable of capturing and integrating sub-word information using character n-gram embeddings

<sup>2</sup>[https://github.com/neerajal504/GenderBias\\_corpus](https://github.com/neerajal504/GenderBias_corpus)

during training (Kunchukuttan et al., 2020).

## 4 Methodology

### 4.1 Bias Statement

Various definitions of bias exist and vary in research as explained in the paper (Blodgett et al., 2020). Our work focuses on stereotypical associations between masculine and feminine gender and professional occupations and emotions in Fast-Text word embeddings. The classic example of "He is a doctor" and "She is a nurse" comes into play here. It is especially harmful to the representation of minority communities, since these stereotypes often end up undermining these communities (Moss-Racusin et al., 2012). Downstream NLP applications learn from these stereotypes, and the risk of discrimination on the basis of gender in this case keeps seeping further into the system.

Our work tries to de-correlate gender with occupation and emotions, which will help reduce bias in these systems.

### 4.2 Quantifying bias for Occupations and Emotions

We use the following methods to quantify the bias before and after debiasing.  $M_{gen}$  is used for gendered attributes like gendered occupations.  $M_{neu}$  is used for neutral attributes like emotions and neutral occupations. We use these two different methods because our data has two different parts — gendered and neutral.

#### 4.2.1 For neutral occupations and emotions:

$M_{neu}$

1. **ECT-n:** Dev and Phillips (2019) use this test to measure bias. We use the target word pairs  $T$ , and the neutral attributes list  $P_{neu}$  and emotions. We separate the target word pairs into masculine and feminine-targeted words respectively. For each of the pairs  $\vec{m}_i, \vec{f}_i$  in  $T$  we create two means  $\vec{a}_1$  and  $\vec{a}_2$ .

$$\vec{a}_1 = \frac{1}{|M|} \sum_{\vec{m} \in M} \vec{m} \quad (1)$$

$$\vec{a}_2 = \frac{1}{|F|} \sum_{\vec{f} \in F} \vec{f} \quad (2)$$

$M$  are masculine word embeddings,  $F$  are feminine word embeddings of the target word pairs  $T$ . We then create two arrays, one with the cosine similarity between the neutral word

embeddings and  $\vec{a}_1$ , the other with the neutral word embeddings and  $\vec{a}_2$ . We calculate the Spearman correlation between the rank orders of these two arrays found. Spearman rank correlation is a non-parametric test that is used to measure the degree of association between two variables. Higher the correlation, the less the bias. The range of the correlation is  $[-1, 1]$ . Ideally, the correlation should be equal to one as the professions or emotions should not depend upon gender. Debiasing should bring the value closer to one.

2. **RND-n:** Relative Norm Distance was first used by Garg et al. (2018). It captures the relative strength of the association of a neutral word with respect to two groups. As shown in equation 3 we average the masculine and feminine-targeted words in  $M, F$  in  $T$  respectively. For every attribute,  $\vec{p}$  in  $P_{neu}$  and emotions we find the norm of the average of the target words and the attribute  $\vec{p}$ . The higher the value of the relative norm, the more biased our professions and emotions are. Debiasing should reduce this value and bring it closer to zero.

$$\sum_{\vec{p} \in P_{neu}} (||avg(M) - \vec{p}||_2 - ||avg(F) - \vec{p}||_2) \quad (3)$$

#### 4.2.2 For gendered occupations: $M_{gen}$

1. **ECT-g:** We use the target word pairs  $T$  and the gendered professions list  $P_{gen}$ . Using  $\vec{a}_1$  found in equation 1 and  $\vec{a}_2$  found in equation 2.  $P_{gen}$  has masculine and feminine profession word pairs. We create two arrays, one with cosine similarity of masculine profession word embeddings and  $\vec{a}_1$ . The other with the cosine similarity of feminine profession word embeddings and  $\vec{a}_2$ . We calculate the Spearman correlation of the rank of these two arrays.

Ideally, there should be a high correlation between these arrays. The masculine profession words' cosine similarity with masculine target words should equal feminine profession words' cosine similarity with feminine target words. The range of the correlation is  $[-1, 1]$ . Higher the correlation, the less the bias. Debiasing should bring the value closer to one.

Emotion	Metrics(%)	Baseline	Projection		Partial Projection	
			PCA	RIPA	PCA	RIPA
Anger	ECT $\uparrow$	57.2	94.5	38.1	<b>99.1</b>	99.0
	RND $\downarrow$	2.0	1.1	0.4	0.7	<b>0.3</b>
Fear	ECT $\uparrow$	86.8	97.8	77.4	<b>99.4</b>	98.9
	RND $\downarrow$	2.6	1.2	0.3	0.8	<b>0.2</b>
Joy	ECT $\uparrow$	75.2	96.1	81.1	<b>99.5</b>	99.2
	RND $\downarrow$	2.5	1.2	0.5	0.8	<b>0.3</b>
Sadness	ECT $\uparrow$	63.1	88.4	56.1	<b>99.4</b>	98.3
	RND $\downarrow$	3.8	1.7	0.7	0.9	<b>0.4</b>

Table 1: Hindi Emotion Results (Principal Component Analysis (PCA), Relational Inner Product Association (RIPA))

Emotion	Metrics(%)	Baseline	Projection		Partial Projection	
			PCA	RIPA	PCA	RIPA
Anger	ECT $\uparrow$	37.9	58.2	52.7	<b>96.1</b>	93.9
	RND $\downarrow$	0.61	0.53	0.10	0.33	<b>0.09</b>
Fear	ECT $\uparrow$	72.6	74.0	71.2	<b>96.4</b>	93.6
	RND $\downarrow$	0.50	0.41	0.11	0.29	<b>0.08</b>
Joy	ECT $\uparrow$	57.3	76.2	58.2	<b>93.4</b>	92.5
	RND $\downarrow$	0.60	0.61	0.13	0.39	<b>0.11</b>
Sadness	ECT $\uparrow$	69.6	80.2	68.9	<b>99.1</b>	96.9
	RND $\downarrow$	0.42	0.37	0.08	0.25	<b>0.07</b>

Table 2: Marathi Emotion Results

- RND-g**: As shown in equation 4 we average the masculine and feminine-targeted words in  $M, F$  in  $T$ , respectively. For every attribute pair  $\vec{p}^1$  and  $\vec{p}^2$  in  $P_{gen}$  we find the norm of the average of the masculine target words and  $\vec{p}^1$ , feminine target words and  $\vec{p}^2$ . The higher the value of the relative norm, the more biased the professions are. Debiasing should reduce this value and bring it closer to zero.

$$\sum_{\vec{p}^1, \vec{p}^2 \in P} (||avg(M) - \vec{p}^1||_2 - ||avg(F) - \vec{p}^2||_2) \quad (4)$$

### 4.3 Debiasing techniques

#### 4.3.1 Finding out the gender subspace

We need a vector  $\vec{v}_b$  that represents the gender direction. We find this in the following ways: using RIPA and PCA.

- RIPA**: [Ethayarajh et al. \(2019\)](#) first used this subspace to capture gender information. We define a bias vector  $\vec{v}_b$  which defines the gender direction. Given the target set  $T$  containing masculine and feminine words, for each

$T_j$  in  $T$ , we find out  $T_f - T_m$  and stack them to create an array.  $T_f$  is the feminine word embedding,  $T_m$  is the masculine word embedding. We find the first principal component using Principal Component Analysis (PCA) of the array found above. This component captures the gender information of the given embeddings.

- PCA**: In this method, given  $T$ , we find out the average  $a$  of the masculine and feminine word embeddings for each given pair. We then compute  $T_f - a$  and  $T_m - a$  for each  $T_j$  in  $T$ . We stack them into an array and find out the first component using the PCA of the above array.

#### 4.3.2 Debiasing methods

[Bolukbasi et al. \(2016\)](#) used Hard Debiasing to mitigate bias. This method needs additional seed words which are then trained on a SVM to find out if the word is biased or neutral. Thus, not all words are debiased in the vocabulary except the chosen ones. This makes the method not fully automatic.

Here we use more straightforward methods to



Metric (%)	Base	Projection		Partial Proj.	
		PCA	RIPA	PCA	RIPA
ECT-n $\uparrow$	86.0	95.6	81.4	<b>99.7</b>	98.9
RND-n $\downarrow$	40.4	19.3	6.2	10.9	<b>6.2</b>
ECT-g $\uparrow$	69	71.4	85.7	<b>90.4</b>	88.0
RND-g $\downarrow$	1.79	1.81	1.79	<b>1.52</b>	1.70

Table 3: Results for Hindi occupations. RND-g is not in % for better readability

Metric (%)	Base	Projection		Partial Proj	
		PCA	RIPA	PCA	RIPA
ECT-n $\uparrow$	51.4	60.6	53.6	<b>99.7</b>	96.5
RND-n $\downarrow$	3.1	3.0	3.0	01.8	<b>0.5</b>
ECT-g $\uparrow$	42.5	21.5	23.8	64.2	<b>76.2</b>
RND-g $\downarrow$	2.58	2.50	2.54	2.02	<b>1.97</b>

Table 4: Results for Marathi occupations. RND-g is not in % for better readability

debias our data.

1. **Projection:** One way to remove bias is to make all the vectors orthogonal to the gender direction. Therefore, we remove the component of  $\vec{v}_b$  from all the vectors. This ensures that there is no component along the gender direction.

$$\vec{w} = \vec{w} - (\vec{w} \cdot \vec{v}_b) \vec{v}_b \quad (5)$$

2. **Partial Projection:** One problem with the debiasing using linear projection is that it changes some word vectors which are gendered by definition, e.g., king, queen. Let  $\mu = \frac{1}{m} \sum_{j=1}^m \mu_j$  where  $\mu_j = \frac{1}{|T_j|} \sum_{t \in T_j} t$  be the mean of a target pair. Here m is the length of T. We suggest the new vector as shown in equation 6. This is similar to the linear projection approach, but instead of zero magnitude along the gender direction, we project a magnitude of constant  $\mu$  along with it. This adds a constant to the debiasing term.

$$\vec{w} = \vec{w} - (\vec{w} \cdot \vec{v}_b) \vec{v}_b + \mu \quad (6)$$

## 5 Results and Discussion

Table 1 and 2 show results for the emotions in Hindi and Marathi respectively. We observe that anger is the most biased in both languages according to the ECT metric as it has the lowest value. Amongst the debiasing techniques, we see that partial projection

with RIPA works the best for the ECT metric and partial projection with PCA works the best for the RND metric.

ECT-n and RND-n are results for neutral occupations, and ECT-g and RND-g are for gendered occupations. Table 3 shows the results for both gendered and neutral occupations in Hindi. We see that partial projection We see that for neutral occupations, partial projection with PCA works the best for ECT and partial projection with RIPA works the best for RND. For gendered occupations, we see that partial projection with PCA works the best for both ECT and RND.

Table 4 shows the results for both gendered and neutral occupations in Marathi. We see that the best results are obtained for neutral occupations with partial projection with PCA for ECT and partial projection with RIPA for RND. For gendered occupations, we see that we get the best results with partial projection with RIPA for ECT and RND.

However, we observe some anomalies in the results when projection debiasing method is used. We hypothesize that completely removing the gender information changes some vectors, which are masculine or feminine, by the grammatical definition of the gender. For example, words like king, grandfather and boy which are masculine by the grammatical definition of gender should preserve their gender information. Hence we note that partial projection performs the best because it has a gender component to it.

## 6 Conclusion and Future work

In this paper, we attempted to find gender stereotypes on occupations and emotions and tried to debias them. Embedding Coherence Test and Relative Norm Distance were used as a bias metric in the gender subspace. The debiasing methods used were projection and partial projection. But we see that partial projection as a debiasing method works the best in most cases.

Future work could include trying out these techniques on downstream tasks and checking the performance before and after debiasing. The main problem with experimenting on downstream tasks is the availability of datasets in these languages. We would also like to experiment with debiasing contextual embeddings and large language models. Apart from that we would also like to address other types of bias like religion, social and cultural, which are particularly inherent in Hindi and

Marathi.

## 7 Acknowledgements

We are grateful to Anna Currey, Krithika Ramesh, Gauri Gupta, Sahil Khose, Soham Tiwari for reviewing the manuscript and helping in writing the bias statement.

## References

- 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.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomáš Mikolov. 2016. [Enriching word vectors with subword information](#). *CoRR*, abs/1607.04606.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29:4349–4357.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Bernard Comrie. 1999. Grammatical gender systems: a linguist’s assessment. *Journal of Psycholinguistic research*, 28(5):457–466.
- Sunipa Dev and Jeff Phillips. 2019. Attenuating bias in word vectors. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 879–887. PMLR.
- Kawin Ethayarajh, David Duvenaud, and Graeme Hirst. 2019. [Understanding undesirable word embedding associations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1696–1705, Florence, Italy. Association for Computational Linguistics.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Svetlana Kiritchenko and Saif M Mohammad. 2018. Examining gender and race bias in two hundred sentiment analysis systems. *arXiv preprint arXiv:1805.04508*.
- Anoop Kunchukuttan, Divyanshu Kakwani, Satish Golla, Avik Bhattacharyya, Mitesh M Khapra, Pratyush Kumar, et al. 2020. Ai4bharat-indicnlp corpus: Monolingual corpora and word embeddings for indic languages. *arXiv preprint arXiv:2005.00085*.
- Vijit Malik, Sunipa Dev, Akihiro Nishi, Nanyun Peng, and Kai-Wei Chang. 2021. Socially aware bias measurements for hindi language representations. *arXiv preprint arXiv:2110.07871*.
- Corinne A Moss-Racusin, John F Dovidio, Victoria L Brescoll, Mark J Graham, and Jo Handelsman. 2012. Science faculty’s subtle gender biases favor male students. *Proceedings of the national academy of sciences*, 109(41):16474–16479.
- Arun K Pujari, Ansh Mittal, Anshuman Padhi, Anshul Jain, Mukesh Jadon, and Vikas Kumar. 2019. Debiasing gender biased hindi words with word-embedding. In *Proceedings of the 2019 2nd International Conference on Algorithms, Computing and Artificial Intelligence*, pages 450–456.
- Krithika Ramesh, Gauri Gupta, and Sanjay Singh. 2021. Evaluating gender bias in hindi-english machine translation. In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 16–23.
- Karolina Stanczak and Isabelle Augenstein. 2021. A survey on gender bias in natural language processing. *arXiv preprint arXiv:2112.14168*.
- 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](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1630–1640, Florence, Italy. Association for Computational Linguistics.
- Pei Zhou, Weijia Shi, Jieyu Zhao, Kuan-Hao Huang, Muhao Chen, Ryan Cotterell, and Kai-Wei Chang. 2019. [Examining gender bias in languages with grammatical gender](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5276–5284, Hong Kong, China. Association for Computational Linguistics.