An Unfair Affinity Toward Fairness: Characterizing 70 Years of Social Biases in B^{H}ollywood

Kunal Khadilkar[1] **and Ashiqur R. KhudaBukhsh**[1]

[1]Carnegie Mellon University

{kkhadilk, akhudabu}@andrew.cmu.edu

Abstract

Bollywood, aka the Mumbai film industry, is one of the biggest movie industries in the world. With a current movie market share of worth 2.1 billion dollars and a target audience base of 1.2 billion people, Bollywood is a formidable entertainment force. While the entertainment impact in terms of lives that Bollywood can potentially touch is mammoth, no NLP study on social biases in Bollywood content exists. We thus seek to understand social biases in a developing country through the lens of popular movies. Our argument is simple – popular movie content reflects social norms and beliefs in some form or shape. We present our preliminary findings on a longitudinal corpus of English subtitles of popular Bollywood movies focusing on (1) social bias toward a fair skin color (2) gender biases, and (3) gender representation. We contrast our findings with a similar corpus of Hollywood movies.

1 Introduction

On a diachronic corpus of popular entertainment, what types of social biases can we analyze and detect? In this paper, we analyze a curated corpus of film subtitles from the Bollywood and Hollywood film industry, spread over 70 years (1950 to 2020) and present our preliminary findings on three aspects (1) social bias toward a fair skin color (2) gender biases, and (3) gender representation. While studies analyzing gender stereotypes across different languages [Lewis and Lupyan, 2020] and detecting bias in word embeddings [Garg *et al.*, 2018] exist, barring few lines of work focusing on a single Bollywood movie [Chatterjee, 2016] or a small subset of movies [Khan and Taylor, 2018], no existing work has focused on analyzing social biases and under-representation in Bollywood films. We perform a first ever large-scale study comprising 1400 films, to uncover implicit social and racial biases in the entertainment industry, using computational science and natural language processing techniques.

In Figure 1, what do the prominent presence of negative adjectives such as `wanton` in old Bollywood movies (see, Figure 4(a)) juxtaposed with the presence of positive verbs such as `respect` (see, Figure 4(b)) tell us? In this paper, we explore a wide array of NLP techniques to analyze our research questions through the lens of popular movies. We use

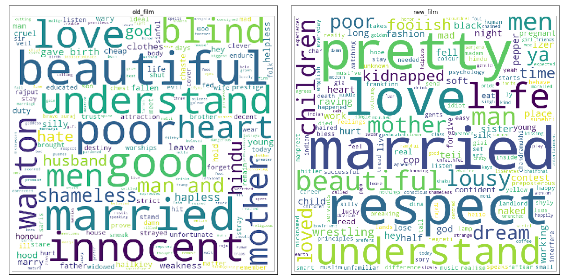

(a) Old Bollywood (b) Recent Bollywood

Figure 1: Word cloud visualizations for the five word contexts containing utterances of 'woman/women'.

1. Simple *count-based statistics* relying on highly popular lexicons [Warriner *et al.*, 2013] and gender representation studies [Twenge *et al.*, 2012; Sendén *et al.*, 2015];
2. *Cloze test*, an analysis technique that have a solid grounding in psycholinguistics literature [Taylor, 1953; Smith and Levy, 2011]. To the best of our knowledge, for the first time, we explore a recent technique [Petroni *et al.*, 2019] previously used to mine political insights [Palakodety *et al.*, 2020] in the context of uncovering social biases. Through a series of cloze tests on a language model [Devlin *et al.*, 2019] fine-tuned on our data sets, we present our findings.
3. *Analysis of aligned diachronic word embedding spaces* using recently proposed techniques [Hamilton *et al.*, 2016].

2 Data Set

We construct the following two data sets.

1. Bollywood movies, $\mathcal{D}_{bolly}$**:** We consider movies spanning seven decades (1950–2020). For each decade, we retrieved subtitles [Lison and Tiedemann, 2016] of 100 movies (700 total films).

2. Hollywood movies, $\mathcal{D}_{holly}$**:** Similar to Bollywood movies, we considered 100 top-grossing movies from each of the seven decades (700 total films). Overall, $\mathcal{D}_{bolly}$ and $\mathcal{D}_{holly}$ consist of 1.1M dialogs (6.2M Tokens) and 1M dialogs (5.4M Tokens), respectively.

Probe	**BERT$_{base}$**	**BERT**$_{\mathcal{D}^{old}_{bolly}}$	**BERT**$_{\mathcal{D}^{new}_{bolly}}$	**BERT**$_{\mathcal{D}^{old}_{holly}}$	**BERT**$_{\mathcal{D}^{new}_{holly}}$
$cloze_1$	soft, beautiful, pale, tanned, smooth	fair, no, pale, tanned, tan	fair, tanned, golden, smooth, pale	fair, pale, blue, golden, gold	fair, pale, tanned, golden, dark
$cloze_2$	woman, man, widow, doctor, slave, soldier, bachelor, merchant, farmer, lawyer, servant **[4.8]**	prostitute, servant, woman, slave, bachelor, doctor, lawyer, man, widow, maid, worker **[4.64]**	doctor, woman, servant, lawyer, maid, man, nurse, teacher, gardener, lady, hindu **[5.7]**	woman, slave, servant, nurse, lady, man, teacher, lawyer, peasant, maid, wife **[5.3]**	woman, lawyer, doctor, nurse, teacher, man, writer, secretary, prostitute, professional, carpenter **[5.7]**
$cloze_3$	man, soldier, gentleman, farmer, merchant, woman, slave, bachelor, doctor, carpenter, servant **[5.48]**	man, gentleman, bachelor, lawyer, servant, doctor, farmer, worker, craftsman, slave, criminal **[5.0]**	doctor, lawyer, policeman, man, farmer, bachelor, gardener, servant, soldier, mechanic, builder **[5.3]**	carpenter, policeman, lawyer, soldier, farmer, gentleman, servant, man, peasant, slave, doctor **[5.0]**	man, lawyer, soldier, doctor, carpenter, gentleman, farmer, clergyman, writer, craftsman, minister **[5.78]**

Table 1: Cloze test results. Predicted tokens are ranked by decreasing probability. Positive and negative words are color coded with blue and red, respectively. The number in the bracket represents the average valence score (obtained from [Warriner *et al.*, 2013]) calculated for the answers to the cloze test.

For a subset of our analyses, we divide our corpus into three buckets: (1) Films from 1950 to 1969 ($\mathcal{D}^{old}_{bolly}$ and $\mathcal{D}^{old}_{holly}$); (2) Films from 1970 to 1999 ($\mathcal{D}^{mid}_{bolly}$ and $\mathcal{D}^{mid}_{holly}$); and (3) Films from 2000 to 2019 ($\mathcal{D}^{new}_{bolly}$ and $\mathcal{D}^{new}_{holly}$).

3 Our Analyses

For our first two analyses of gendered pronouns and cloze tests, we restrict ourselves to the temporally extreme subsets of old and new movie sub-corpora.

3.1 Pronouns as a Proxy for Representation:

Following extensive literature on gendered pronouns' relative distributions and their implications [Twenge *et al.*, 2012; Sendén *et al.*, 2015], we first consider a simple measure of gender representation: relative occurrence of pronouns of each gender (Men: he, him. Women: she, her). Let $\mathcal{N}_w$ denote the number of times a token w appears in a corpus. We define *Male Pronoun Ratio (MPR)* as follows:
$MPR = \frac{\mathcal{N}_{he}+\mathcal{N}_{him}}{\mathcal{N}_{he}+\mathcal{N}_{him}+\mathcal{N}_{she}+\mathcal{N}_{her}} * 100.$
Figure 2 plots *MPR* of our decade-wise movie data sets and contrasts with *MPR* computed using google n-grams. While compared to 1950s both Bollywood and Hollywood demonstrate considerable decrease in MPR, our results indicate that even now, both Bollywood and Hollywood exhibit comparable skew in gendered pronoun usage.

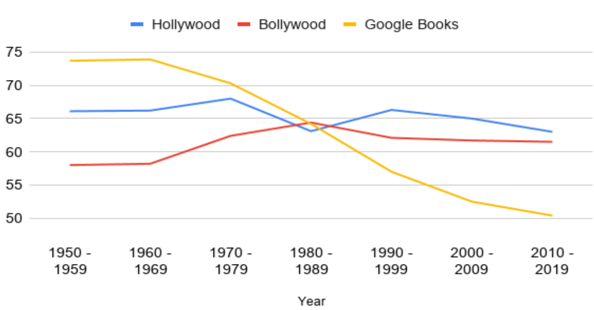

Figure 2: MPR in $\mathcal{D}_{bolly}$ and $\mathcal{D}_{holly}$

3.2 Cloze Test Using Language Models:

When presented with a sentence (or a sentence stem) with a missing word, a cloze task [Taylor, 1953] is essentially a fill-in-the-blank task. For instance, in the following cloze task: *In the* `[MASK]`*, it is very sunny*, `summer` is a likely completion for the missing word. Given a cloze test, `BERT`, a well-known language model [Devlin *et al.*, 2019], outputs a series of token ranked by probability. In fact, in the above cloze test, the top three tokens (ranked by probability) predicted by $\texttt{BERT}_{base}$ are: summer, winter and spring. Recent lines of research has explored `BERT`'s masked query prediction for (1) knowledge base extraction [Petroni *et al.*, 2019] and (2) mining political insights [Palakodety *et al.*, 2020].

Following [Palakodety *et al.*, 2020], we fine-tune `BERT` on four sub-corpora: $\mathcal{D}^{old}_{bolly}$, $\mathcal{D}^{old}_{holly}$,$\mathcal{D}^{old}_{bolly}$, and $\mathcal{D}^{old}_{holly}$, with 100 movies in each corpora. We denote the pretrained `BERT` model as $\texttt{BERT}_{base}$ and a fine-tuned `BERT` on corpus $\mathcal{D}$ as $\texttt{BERT}_{\mathcal{D}}$. Out of a thorough analysis with several cloze tests with phrase variations, in Table 1 we report three such results for the following cloze tests: (1) *A beautiful woman should have* `[MASK]` *skin*. ($cloze_1$); (2) *A woman should be a* `[MASK]` *by occupation*. ($cloze_2$); and (3) *A man should be a* `[MASK]` *by occupation* ($cloze_3$).

For all cloze tests, we evaluate robustness to phrase variations to ensure that our results of the masked words are qualitatively similar. For instance, one of the phrase variations of $cloze_2$ we consider is *A woman should be a* `[MASK]` *by profession*. Similarly, one of the phrase variations for $cloze_1$ we consider is *For a woman to be beautiful, she should have* `[MASK]` *skin*. Overall, our findings remain qualitatively similar across the phrase variations.

Summary of findings
Given cultural and social bias toward fair skin prevalent in India for years [Karan, 2008], it is not surprising that Bollywood exhibits an affinity toward fairness. However, our analysis indicates Hollywood also shows bias toward lighter skin color. Further, we observe cloze completions for both genders across both movie industries improve over time. As compared to men, the completions for women in older Bollywood movies exhibit stronger negative terms such as `slave` and `prostitute`.

	Bollywood	Hollywood
Women	22%	7%
Men	5%	15%

Table 2: Percent increase in average valence score for cloze test completions between old movies and new movies.

We now present a quantitative evaluation of our completion results. We consider a well-known lexicon of emotional valence ratings [Warriner *et al.*, 2013] of nearly 14,000 English words to quantify the change of cloze test completions over time. The valence score of these words are presented in a scale of 1 to 10 with 10 indicating highly positive and 1 indicating highly negative. For example, the emotional valence score of `happy` and `sad` are 8.47 and 2.10, respectively. For a given data set and a cloze test pair, we compute the average valence score of the completions (listed in square brackets in Table 1). We further note that comparing completion words across the two genders can be technically difficult simply because we observe that there can be potential bias in the scores. For instance, the valence scores for `man` and `woman` are 5.42 and 7.09, respectively. Hence we restrict ourselves to comparing within a specific gender for a given movie industry.

Table 2 lists the percentage of increase in the valence score of the completion for a particular gender across different movie industries. We note that for both Bollywood and Hollywood, the valence scores for both genders improved over time. However, for Bollywood, we notice that rate of increase for women is substantially more pronounced than that for men. This observation aligns with the continual fight for gender equality in India [Nayak and Mahanta, 2012] and major movements that have mobilised voices for women's right to work [Chen, 1995], financial independence [Goyal and Parkash, 2011], and marital laws [Nigam, 2005].

3.3 Aligning Diachronic Word Embeddings

The meaning of words and the context in which they are used change over time [Xie *et al.*, 2019]. The language spoken in a community is representative of the cultural norms and customs followed in that region. Existing hypotheses [Bybee, 2006] indicate that word frequency may play a role in changing the meaning of the words over time. The meanings of less-frequent words are more susceptible to drifts than that of highly frequent words. [Hamilton *et al.*, 2016] provide a robust multilingual approach to align diachronic word embeddings using orthogonal procrustes. We follow the same method to align different sub-corpora for Bollywood and Hollywood.

We train word2vec [Mikolov *et al.*, 2013] with SGNS (Skip-gram with Negative Sampling), to create embeddings for each of the buckets mentioned above. Let $\mathcal{W}^{(t)} \in \mathbb{R}^{(d) \times \mathcal{V}}$ be the matrix of word embeddings learnt for period t for vocabulary $\mathcal{V}$. Following [Hamilton *et al.*, 2016], we align the word embeddings using the top 10,000 common tokens present across time periods t and $t+1$ by optimizing:

$$R^t = arg \min_{Q^T Q = I} \left\| Q\mathcal{W}^t - \mathcal{W}^{(t+1)} \right\|_F \quad (1)$$

where $R^t \in \mathbb{R}^{d \times d}$.

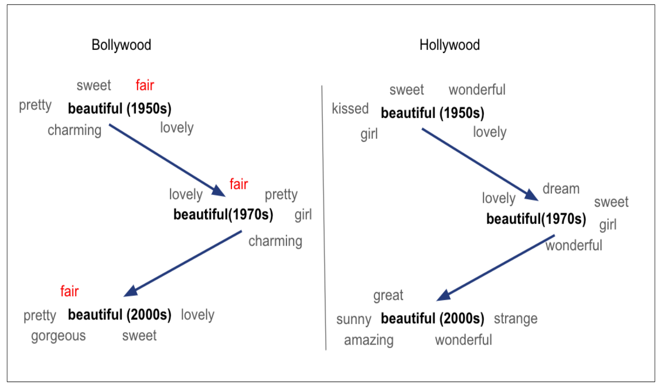

Figure 3: Nearest neighbors of `beautiful` over the years

Affinity toward Fairness

Figure 3 visualizes the nearest neighbors of `beautiful` in our aligned embedding spaces of Hollywood and Bollywood sub-corpora. As shown in Figure 3, the age-old affinity toward lighter skin in Indian culture [Dlova *et al.*, 2015; Chattopadhyay, 2019; Madhukalya, 2020] is reflected through the consistent presence of `fair` among the nearest neighbors of all three Bollywood sub-corpora. Although our cloze tests indicates Hollywood also exhibits bias towards lighter skin color, our diachronic word embedding analysis reveals that possibly the bias is less pronounced than in Bollywood.

Portrayal of Women and Men

We next focus on the portrayal of women and men. We observe that the valence scores for both genders across both movie industries show a similar pattern. The scores are the lowest during the 1970-1999 time period. The valence scores for the newer movies are better than the scores for the older movies. The dip in the valence scores during the time-period of 1970-1999 in India can be ascribed to a social and cultural crisis influenced by an unstable political climate (assassinations of two prime ministers [Hardgrave Jr, 1985; Kaarthikenyan and Raju, 2008]), two major wars between India and Pakistan [Schofield, 2010; Bose, 2009], and a large overlap with a pre-economic liberalization period [Pedersen, 2000].

Social Evil

Our next focus is on a social evil discriminatory toward women. We consider the 'dowry' system that has plagued the Indian society for a long time [Dalmia and Lawrence, 2005]. Dowry refers to a transaction of tangible financial objects in the form of durable goods, cash, and real or movable property between the bride's family gives and the bridegroom, his parents and his relatives as a condition of the marriage Although, legally, dowry has been prohibited in India since 1961 [Rao, 1973], this practice has continued well after its legal prohibition and has a strong link to social crises such as female foeticide [Ghansham, 2002], domestic abuse and violence [Banerjee, 2014; Rastogi and Therly, 2006], and dowry deaths [Ahmad, 2008]. However, while the practice continued abated, recent studies have reported positive changes in the society where the general attitude towards the system has become negative [Srinivasan and Lee, 2004].

Probe	$\mathbf{BERT}_{\mathbf{base}}$	$\mathbf{BERT}_{\mathcal{D}^{old}_{bolly}}$	$\mathbf{BERT}_{\mathcal{D}^{new}_{bolly}}$
$cloze_4$	corruption, poverty, malaria, pollution, hunger, terrorism, unemployment, drought, famine, war, tourism	poverty, love, war, hunger, unemployment, india, famine, money, marriage, education, kashmir	poverty, pakistan, kashmir, terrorism, corruption, india, drugs, dowry, unemployment, hunger, rape

Table 3: Cloze test results.

Probe	$\mathbf{BERT}_{\mathbf{base}}$	$\mathbf{BERT}_{\mathcal{D}^{old}_{holly}}$	$\mathbf{BERT}_{\mathcal{D}^{new}_{holly}}$
$cloze_5$	poverty, corruption, unemployment, crime, terrorism, racism, pollution, hunger, war, cancer, inequality	war, poverty, money, unemployment, slavery, immigration, alcoholism, education, imperialism, russia, hunger	poverty, slavery, immigration, unemployment, money, war, racism, hunger, communism, america, education

Table 4: Cloze test results.

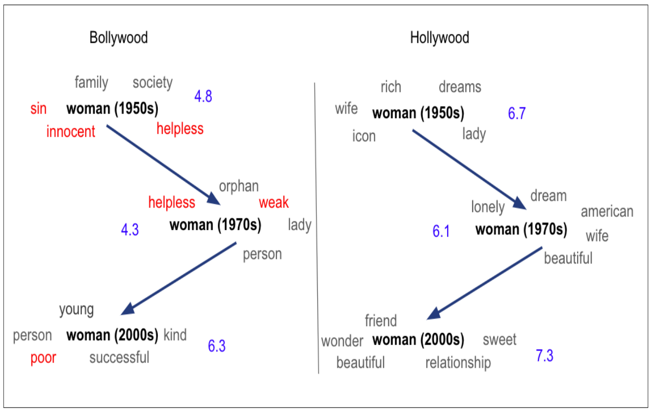

(a) `Woman` over the years

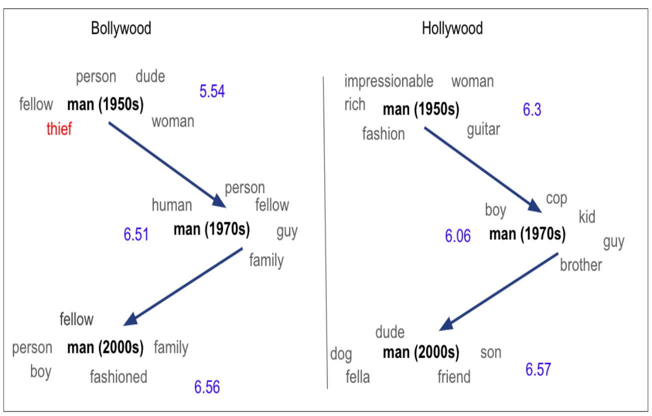

(b) `Man` over the years

Figure 4: Nearest neighbors of `man` and `woman` over the years. The overall average valence of nearest neighbors according to the lexicon provided in [Warriner *et al.*, 2013] for a given time-period is presented in blue font.

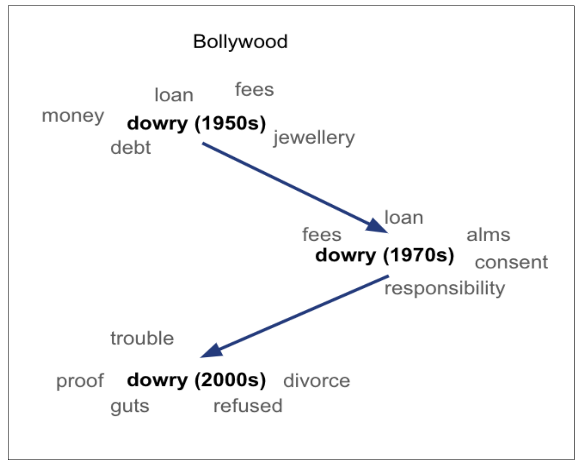

Figure 5: Nearest neighbors of `dowry` over the years.

We were curious to analyze how word embedding spaces reflect the attitude toward `dowry`. As shown in Figure 5, we observe that while nouns such as `money`, `debt`, `jewellery`, `fees`, and `loan` are the nearest neighbors in older films indicating compliance to this practice, modern films exhibit non-compliance (e.g., `guts` and `refused`) and indicate some of the consequences of such non-compliance (e.g., `divorce` and `trouble`) in the form of nearest neighbors. We find that indeed, films provide a snapshot of cultural values of a particular country thus allowing us to gauge the progress of a nation over time.

4 Future Directions

Our future lines of work include (1) tracking moral sentiment [Xie *et al.*, 2019]; (2) exploring debiasing techniques [Manzini *et al.*, 2019]; (3) analyzing representation of other minorities (e.g., Muslims, dalits and the LGBTQ community); and (4) exploring a newly-proposed machine-translation data set distance measure [KhudaBukhsh *et al.*, 2020] to uncover contrasting social aspects.

While our current focus is on characterizing gender bias and representation, our rich data set has untapped potential in analyzing a wide range of research questions. For instance, in line with the cloze tests conducted in [Palakodety *et al.*, 2020], Table 3 and 4 list the results for two additional cloze tests: (1) *The biggest problem of India is* `[MASK]`. (referred as $cloze_4$) - For Bollywood and (2) *The biggest problem of America is* `[MASK]`. (referred as $cloze_5$) - For Hollywood.

We observe that the dynamic political conditions are reflected in the completion results (e.g., `Kashmir`, `Pakistan` and `Russia`) [Schofield, 2010; Bose, 2009; LaFeber and Abbott, 1972]. We also note that the list of ongoing problems in the U.S. contains the major issue on which the 2020 election is being fought over: racism.

References

[Ahmad, 2008] Nehaluddin Ahmad. Dowry deaths (bride burning) in india and abetment of suicide: a socio-legal appraisal. *JE Asia & Int'l L.*, 1:275, 2008.

[Banerjee, 2014] Priya R Banerjee. Dowry in 21st-century india: the sociocultural face of exploitation. *Trauma, Violence, & Abuse*, 15(1):34–40, 2014.

[Bose, 2009] Sumantra Bose. *Kashmir: Roots of conflict, paths to peace*. Harvard University Press, 2009.

[Bybee, 2006] Joan Bybee. *Frequency of Use and the Organization of Language*. 12 2006.

[Chatterjee, 2016] Sushmita Chatterjee. 'english vinglish'and bollywood: what is 'new'about the 'new woman'? *Gender, Place & Culture*, 23(8):1179–1192, 2016.

[Chattopadhyay, 2019] Sriya Chattopadhyay. *Fair-Unfair: Prevalence of Colorism in Indian Matrimonial Ads and Married Women's Perceptions of Skin-Tone Bias in India*. PhD thesis, Bowling Green State University, 2019.

[Chen, 1995] Martha Chen. *A matter of survival: Women's right to employment in India and Bangladesh*, volume 38. Oxford: Clarendon Press, 1995.

[Dalmia and Lawrence, 2005] Sonia Dalmia and Pareena G Lawrence. The institution of dowry in india: Why it continues to prevail. *The Journal of Developing Areas*, pages 71–93, 2005.

[Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, June 2019.

[Dlova *et al.*, 2015] Ncoza C Dlova, Saja H Hamed, Joyce Tsoka-Gwegweni, and Anneke Grobler. Skin lightening practices: an epidemiological study of south african women of african and indian ancestries. *British Journal of Dermatology*, 173:2–9, 2015.

[Garg *et al.*, 2018] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644, 2018.

[Ghansham, 2002] Devaki Monani Ghansham. Female foeticide and the dowry system in india. In *Townsville International Women's Conference, James Cook Univ., Australia*, 2002.

[Goyal and Parkash, 2011] Meenu Goyal and Jai Parkash. Women entrepreneurship in india-problems and prospects. *International journal of multidisciplinary research*, 1(5):195–207, 2011.

[Hamilton *et al.*, 2016] William L. Hamilton, Jure Leskovec, and Dan Jurafsky. Diachronic word embeddings reveal statistical laws of semantic change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany, August 2016. Association for Computational Linguistics.

[Hardgrave Jr, 1985] Robert L Hardgrave Jr. India in 1984: Confrontation, assassination, and succession. *Asian Survey*, 25(2):131–144, 1985.

[Kaarthikenyan and Raju, 2008] DR Kaarthikenyan and Radhavinod Raju. *Rajiv Gandhi Assassination*. Sterling Publishers Pvt. Ltd, 2008.

[Karan, 2008] Kavita Karan. Obsessions with fair skin: Color discourses in indian advertising. *Advertising & society review*, 9(2), 2008.

[Khan and Taylor, 2018] Subuhi Khan and Laramie Taylor. Gender policing in mainstream hindi cinema: A decade of central female characters in top-grossing bollywood movies. *International Journal of Communication*, 12:22, 2018.

[KhudaBukhsh *et al.*, 2020] Ashiqur R. KhudaBukhsh, Rupak Sarkar, Mark S. Kamlet, and Tom M. Mitchell. We don't speak the same language: Interpreting polarization through machine translation, 2020.

[LaFeber and Abbott, 1972] Walter LaFeber and Brian Abbott. *America, Russia, and the Cold War, 1945-1975*. Wiley, 1972.

[Lewis and Lupyan, 2020] Molly Lewis and Gary Lupyan. Gender stereotypes are reflected in the distributional structure of 25 languages. *Nature human behaviour*, pages 1–8, 2020.

[Lison and Tiedemann, 2016] Pierre Lison and Jörg Tiedemann. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles. 2016.

[Madhukalya, 2020] Anwesha Madhukalya. These dialogues from bollywood blockbusters are so sexist that you'll want to pull your hair out, 2020.

[Manzini *et al.*, 2019] Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in word embeddings. In *NAACL-HLT*, pages 615–621, June 2019.

[Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc., 2013.

[Nayak and Mahanta, 2012] Purusottam Nayak and Bidisha Mahanta. Women empowerment in india. *Bulletin of Political Economy*, 5(2):155–183, 2012.

[Nigam, 2005] Shalu Nigam. Understanding justice delivery system from the perspective of women litigants as victims of domestic violence in india (specifically in the context of section 498-a, ipc). *Nigam Shalu (2005) Understanding Justice Delivery System From The Perspective of Women Litigants as Victims of Domestic Violence in India (Specifically in the Context of Section 498-A IPC), Occasional Paper*, (35), 2005.

[Palakodety *et al.*, 2020] Shriphani Palakodety, Ashiqur R. KhudaBukhsh, and Jaime G. Carbonell. Mining insights from large-scale corpora using fine-tuned language models. In *Proceedings of the Twenty-Fourth European Conference on Artificial Intelligence (ECAI-20)*, pages 1890–1897, 2020.

[Pedersen, 2000] Jørgen Dige Pedersen. Explaining economic liberalization in india: state and society perspectives. *World Development*, 28(2):265–282, 2000.

[Petroni *et al.*, 2019] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. Language models as knowledge bases? 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 2463–2473, Hong Kong, China, November 2019. Association for Computational Linguistics.

[Rao, 1973] R Jaganmohan Rao. Dowry system in india—a socio-legal approach to the problem. *Journal of the Indian Law Institute*, 15(4):617–625, 1973.

[Rastogi and Therly, 2006] Mudita Rastogi and Paul Therly. Dowry and its link to violence against women in india: Feminist psychological perspectives. *Trauma, Violence, & Abuse*, 7(1):66–77, 2006.

[Schofield, 2010] Victoria Schofield. *Kashmir in conflict: India, Pakistan and the unending war*. Bloomsbury Publishing, 2010.

[Sendén *et al.*, 2015] Marie Gustafsson Sendén, Sverker Sikström, and Torun Lindholm. "she" and "he" in news media messages: Pronoun use reflects gender biases in semantic contexts. *Sex Roles*, 72(1-2):40–49, 2015.

[Smith and Levy, 2011] Nathaniel Smith and Roger Levy. Cloze but no cigar: The complex relationship between cloze, corpus, and subjective probabilities in language processing. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 33, 2011.

[Srinivasan and Lee, 2004] Padma Srinivasan and Gary R Lee. The dowry system in northern india: Women's attitudes and social change. *Journal of Marriage and Family*, 66(5):1108–1117, 2004.

[Taylor, 1953] Wilson L Taylor. "cloze procedure": A new tool for measuring readability. *Journalism quarterly*, 30(4):415–433, 1953.

[Twenge *et al.*, 2012] Jean M Twenge, W Keith Campbell, and Brittany Gentile. Male and female pronoun use in us books reflects women's status, 1900–2008. *Sex roles*, 67(9-10):488–493, 2012.

[Warriner *et al.*, 2013] Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207, 2013.

[Xie *et al.*, 2019] Jing Yi Xie, Renato Ferreira Pinto Junior, Graeme Hirst, and Yang Xu. Text-based inference of moral sentiment change. 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 4654–4663, November 2019.