

# Unequal Representations: Analyzing Intersectional Biases in Word Embeddings Using Representational Similarity Analysis

Michael A. Lepori

Department of Computer Science

Johns Hopkins University

mlepori19@gmail.com

## Abstract

We present a new approach for detecting human-like social biases in word embeddings using representational similarity analysis. Specifically, we probe contextualized and non-contextualized embeddings for evidence of intersectional biases against Black women. We show that these embeddings represent Black women as simultaneously less feminine than White women, and less Black than Black men. This finding aligns with intersectionality theory, which argues that multiple identity categories (such as race or sex) layer on top of each other in order to create unique modes of discrimination that are not shared by any individual category.

## 1 Introduction

Word embeddings are ubiquitous components of modern natural language processing systems, and have enabled performance increases on a wide variety of NLP tasks (Devlin et al., 2019; Peters et al., 2018; Pennington et al., 2014). They provide vector representations of words, and are trained on a large corpus of text in order to capture the distributional semantics of lexical items. Many studies have shown that these embeddings also encode a variety of dangerous social biases, as these biases are present in naturally-occurring English text (see Mehrabi et al. (2019) for a survey of such work).

The present work studies *intersectional biases* in both contextualized and non-contextualized word embeddings. The theory of intersectionality states that multiple different aspects of a person’s identity often combine to create unique modes of discrimination (Crenshaw, 1989; Crenshaw, 1990). For example, specific biases against Black women are not necessarily shared by either White women or Black men. Furthermore, discourses on discrimination and public health typically focus on sexism *or* racism, each implicitly excluding the other (Crenshaw, 1989; Bowleg, 2012).

It is further argued that these discourses are focused on the most privileged members of a group (Crenshaw, 1989). Thus, we might expect that the ostensibly race-neutral category ‘Woman’ is implicitly White, and that the ostensibly non-gendered category ‘Black’ is implicitly male. Using representational similarity analysis (RSA; Kriegeskorte et al. (2008)), we find that both non-contextualized GloVe embeddings (Pennington et al., 2014) and contextualized BERT embeddings (Devlin et al., 2019) display exactly these biases. We also find that the BERT embeddings of the names of Black women are semantically impoverished, compared to the embeddings of the names of White women.

**Related Work:** Many recent studies have sought to understand social biases in NLP systems. Notably, Caliskan et al. (2017) showed that the human biases revealed through the Implicit Association Test (Greenwald et al., 1998) are also present in word embeddings, prompting other studies that seek to detect bias in word embeddings (Kurita et al., 2019; Garg et al., 2018). Other studies have attempted to detect bias in other NLP systems (Rudinger et al., 2018), to automatically discover and measure intersectional biases in word embeddings (Guo and Caliskan, 2020), to analyze the distributions of gendered words in training sets (Zhao et al., 2019; Rudinger et al., 2017), and to automatically mitigate the effects of social biases on NLP systems (Sun et al., 2019). Another study has produced empirical evidence for the claims of intersectionality theory by leveraging distributional representations (Herbelot et al., 2012).

---

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

Perhaps most relevant, some prior work has attempted to understand whether intersectional differences are present in contextualized word embeddings by trying to detect whether these embeddings encoded intersectional stereotypes (May et al., 2019; Tan and Celis, 2019). However, detecting these stereotypes is only a proxy for detecting intersectional biases in embeddings: An embedding might encode intersectional differences without encoding more complex intersectional stereotypes. This work presents a method that leverages representational similarity analysis to directly probe word embeddings for intersectional biases.<sup>1</sup>

## 2 Methods: Representational Similarity Analysis

Representational similarity analysis (RSA) is a technique first used in cognitive neuroscience for analyzing distributed activity patterns in the brain (Kriegeskorte et al., 2008), but recent work has used it to analyze artificial neural systems (Lepori and McCoy, 2020; Chrupała, 2019; Chrupała and Alishahi, 2019; Abnar et al., 2019; Bouchacourt and Baroni, 2018). RSA analyzes the representational geometry of a system, which is defined by the pairwise dissimilarities between representations of a set of stimuli. Specifically, we use RSA to compare the representational geometry of the word embeddings under study to the representational geometries of well-understood *hypothesis models*. If the representational geometries are similar, you can infer that the word embeddings represent the particular information expressed by a hypothesis model.

**Applying RSA:** The approach proceeds as follows: First, we create a corpus that consists of three types of items: group 1 items,  $G_1$ , group 2 items,  $G_2$ , and concept items,  $C$ . Next, we define a reference model  $M_{Ref}$ , which consists of the set of embeddings that we are interested in. Then, we define two hypothesis models,  $M_{Hyp1}$  and  $M_{Hyp2}$ .  $M_{Hyp1}$  instantiates the hypothesis that: ‘Group 1 is associated with the concept under study, while Group 2 is not’. Likewise  $M_{Hyp2}$  instantiates the hypothesis that ‘Group 2 is associated with the concept under study, while Group 1 is not’. We then draw a sample  $c$ , consisting of  $n_1$  items from  $G_1$ ,  $n_2$  items from  $G_2$ , and  $n_3$  items from  $C$ , and calculate the  $(n_1 + n_2 + n_3) \times (n_1 + n_2 + n_3)$  representational geometries  $R$  of each model using a dissimilarity metric  $D$ . In all analyses,  $n_1 = n_2 = n_3 = 10$ .

$$R_{Ref} = D(M_{Ref}, c) \quad (1)$$

$$R_{Hyp1} = D(M_{Hyp1}, c) \quad (2)$$

$$R_{Hyp2} = D(M_{Hyp2}, c) \quad (3)$$

We define  $D = 1 - Spearman's \rho$  for  $M_{Ref}$ . For  $M_{Hyp1}$ , the dissimilarity metric is equal to 1 if one item is from  $G_2$  and the other item is from  $G_1$  or  $C$ , and 0 otherwise. Likewise, the dissimilarity metric for  $M_{Hyp2}$  is equal to 1 if one item is from  $G_1$  and the other item is from  $G_2$  or  $C$ , and 0 otherwise. Clearly, these hypothesis models are not expected to provide a very good fit to the representational geometry of a set of word embeddings. However, if one hypothesis model exhibits a significantly *better* fit than the other, then there is evidence for a problematic bias.

Finally, we calculate the similarity,  $s$ , between the representational geometries of our hypothesis models and our reference model using a similarity metric,  $sim$ . Because the  $R$  matrices are symmetric,  $sim$  only operates on the upper triangle of each  $R$  matrix. We define  $sim = Spearman's \rho$ .

$$s_{Hyp1} = sim(R_{Ref}, R_{Hyp1}) \quad (4)$$

$$s_{Hyp2} = sim(R_{Ref}, R_{Hyp2}) \quad (5)$$

We then repeat the process on 100 samples from our corpus in order to create two 100-length vectors of representational similarities,  $S_{Hyp1}$  and  $S_{Hyp2}$ . We can then apply a nonparametric sign test to the difference of these vectors,  $S_{Hyp1} - S_{Hyp2}$ , in order to test for a consistent difference between measurements of  $s_{Hyp1}$  and  $s_{Hyp2}$ . See Figure 1 for a visual summary of this approach.

<sup>1</sup>Code and data found at [https://github.com/mlepori1/Unequal\\_Representations](https://github.com/mlepori1/Unequal_Representations).

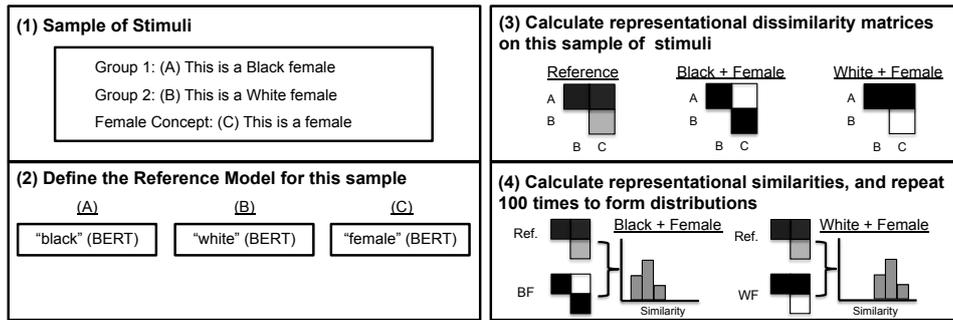


Figure 1: A summary of our approach, applied to BERT embeddings.

### 3 Data

In order to perform the analyses described in Section 2, we need to create sets of group items for all of the groups under study, as well as concept items for the identity concepts under study. For analyses involving GloVe embeddings, these glossaries can only contain single-word entries. For analyses involving BERT embeddings, these entries can be full sentences. See Appendix A for more details about all of the following datasets.

#### 3.1 Single-Word Sets

**Group Items:** For our single-word group items for the Black women, Black men, White women, and White men groups, we use the dataset curated in Sweeney (2013). This dataset includes names commonly associated with all of these identity groups. We rely on name data because identity categories like *Black woman* do not have single word terms (May et al., 2019). See Section 5 for a discussion of the limitations of using name data.

**Concept Items:** For our single-word, ostensibly race-neutral concept items that are associated with the female identity category, we combine the datasets from Nosek et al. (2002a) and Nosek et al. (2002b). The United States Census Bureau defines the Black/African American identity category as “all individuals who identify with one or more nationalities or ethnic groups originating in any of the black racial groups of Africa” (Census Bureau, 2020). They proceed to non-exhaustively list several nationalities that meet this definition. All 10 single-word country names that correspond to the nationalities listed are included in our set of single-word, ostensibly gender-neutral concept items that are associated with the Black identity category. We also include the word *Black*, as well as the word *Africa*, in keeping with the Census Bureau’s definition. See Section 5 for a discussion of the limitations of this approach.

#### 3.2 Multi-Word Sets

We leverage the semantically bleached sentence templates introduced by May et al. (2019) in order to generate our sentence data. These sentences “make heavy use of deixis and are designed to convey little specific meaning beyond that of the terms inserted into them”. Some examples include: “This is a <word/phrase>” and “The <word/phrase> is here”.

**Group Items:** For the Black female group, we fill in the semantically-bleached templates with either the word *Black* concatenated with either *woman*, *women*, *female*, *females*, *girl*, or *girls* (such that the template sentence exhibits proper number agreement). For example, the dataset contains the sentence: *The Black woman is here*. Similarly for the White female group. These sentences are also used by May et al. (2019). Finally, for the Black male group, we fill in the template with the word *Black* and either *man*, *men*, *male*, *males*, *boy*, or *boys*.

**Concept Items:** Our female concept items are chosen to be ostensibly race-neutral. Thus, we fill in the template sentences with one of the following words: *woman*, *women*, *female*, *females*, *girl*, or *girls*. Our Black concept items are chosen to be ostensibly gender-neutral. Thus we fill in the template sentences with either *Black person* or *Black people*.

## 4 Experiments and Results

### 4.1 GloVe Experiments

GloVe embeddings are non-contextualized word embeddings, and so they represent information from one word, ignoring the specific context in which they are found. These experiments seek to test whether RSA can be used to identify intersectional biases in GloVe embeddings. We use 300-dimensional embeddings trained on the Wikipedia 2014 + Gigaword 5 corpus throughout this study.

**Black Female Name Embedding Check** First, we ensure that the GloVe embeddings of the names of Black women contain sufficient semantic detail, such that more in-depth experiments can be performed. Our concept items are a set of words associated with the female identity category. Our Group 1 items are a set of proper names associated with Black women, and our Group 2 items are a set of proper names associated with Black men. See Section 3 for details about these sets. We would expect that the reference model will display greater representational similarity to the hypothesis model that associates female attributes to Black female names. From Table 1 we see that this happens.

**Female Concept** We investigate whether the GloVe embeddings of the names of Black women encode the ‘female’ concept less than the embeddings of the names of White women. Our concept items are a set of words associated with the female identity category. Our Group 1 items are a set of proper names associated with Black women, and our Group 2 items are a set of proper names associated with White women. See Section 3 for details about these sets. From Table 1, we see that the reference model exhibits greater representational similarity to the hypothesis model that associates the concept items to White female names.

**Black Concept** We investigate whether the GloVe embeddings of the names of Black women encode the ‘Black’ concept less than the embeddings of the names of Black men. Our concept items are a set of words associated with the Black identity category. Our Group 1 items are a set of proper names associated with Black women, and our Group 2 items are a set of proper names associated with Black men. From Table 1, we see that the reference model exhibits greater representational similarity to the hypothesis model that associates the concept items with Black male names.

### 4.2 BERT Experiments

BERT embeddings are contextualized word embeddings, and so they represent information from a word in context. This allows us to analyze group terms, using the semantically-bleached sentence data from May et al. (2019), which are designed to convey little meaning beyond that of the group term that we insert into them. We also attempt to reproduce our findings using the single-word proper name data used in Section 4.1. We use the pretrained BERT-base-uncased model for all multi-word studies.

**Black Female Name Embedding Check:** First, we use the pretrained BERT-base-cased model to attempt to reproduce our results from Section 4.1. However, we see from Table 1 that the reference model displays greater representational similarity to the hypothesis model that associates the female concept items to Black *male* names. This indicates that the BERT’s embeddings for Black female names are quite semantically impoverished, and do not reliably encode information about identity categories. Thus, we do not attempt to replicate our investigations of intersectional biases using BERT’s embeddings of proper names. Notably, this effect is *not* reproduced when analyzing White names, and thus demonstrates a problematic asymmetry between the BERT representations of White and Black female names.

**Female Concept:** We investigate whether the BERT embeddings of the word *Black* encode the ‘female’ concept less than the embeddings of the word *White*. Our concept items are a set of sentences containing words associated with the female identity category. Our Group 1 items are as a set of sentences referencing Black women, our Group 2 items are a set of sentences referencing White women. For each sentence in the Group sets, we extract the BERT embedding of the word *Black* or *White*, and for each sentence in the Concept set, we extract the embedding of the word associated with the female identity category (*woman*, *female*, etc.). See Section 3 for more details. From Table 1, we see that the

reference model exhibits greater representational similarity to the hypothesis model that associates the female concept items with White women.

**Black Concept:** We investigate whether the BERT embeddings of words associated with the female identity category encode the ‘Black’ concept less than the embeddings of words associated with the male identity category. Our concept items are a set of sentences containing the phrase *Black person* or *Black people*. Our Group 1 items are a set of sentences referencing Black women, our Group 2 items are a set of sentences referencing Black men. For each sentence in the Group sets, we extract the BERT embedding of the word associated with either the female or male identity category, and for each sentence in the Concept set, we extract the embedding of the word *Black*. See Section 3 for more details. From Table 1, we see that the reference model exhibits greater representational similarity to the hypothesis model that associates the Black concept items with Black men.

Embedding	Group 1	Group 2	Concept	Group 1 + Con	Group 2 + Con
GloVe	B F Names	B M Names	Female	<b>.373</b>	.189
GloVe	B F Names	W F Names	Female	.093	<b>.666</b>
GloVe	B F Names	B M Names	Black	.174	<b>.298</b>
BERT	B F Names	B M Names	Female	.119	<b>.307</b>
BERT	W F Names	W M Names	Female	<b>.368</b>	.021
BERT	B F Sent	W F Sent	Female	.050	<b>.110</b>
BERT	B F Sent	B M Sent	Black	.253	<b>.289</b>

Table 1: Results for all experiments. All differences statistically significant ( $p < .001$ )

## 5 Limitations

One notable limitation of this work is that we study artificial binary gender and race distinctions. Unfortunately, these binary distinctions exclude many individuals from our analysis. Future work should expand upon the methodological framework presented here in order to assess biases against groups that are not represented by these distinctions. Furthermore, Gaddis (2017) describes several limitations of relying on name data to represent racial categories. Perhaps most importantly, that work demonstrates that ‘Black names’ are not one monolithic category, that perceptions of the Blackness of names are correlated with other confounding variables (such as education status), and that the associations between identity categories and proper names are only loosely reflected in real-world naming practices. Furthermore, we rely on an incomplete set of country names provided by the Census Bureau to generate concept sets for the Black identity category. We acknowledge that this excludes many nationalities from our analysis, and also acknowledge that any attempt to exhaustively list nations associated with the Black identity category is destined to fail. We hope that our results inspire future work that aspires to more robust coverage of the Black identity group. Finally, it is important to note that this method can only provide positive evidence: If this method does not detect intersectional biases for a particular system, that does not provide evidence that the system does not contain such biases.

## 6 Conclusion

Both contextualized and non-contextualized word embeddings learn to represent human-like biases. We apply RSA to the task of detecting these biases, and find that both GloVe embeddings and BERT embeddings reproduce (and thus perpetuate) the dual marginalization of Black women within both the Black community and female community, as predicted by the theory of intersectionality. Specifically, we showed that word embeddings represent the female identity category as implicitly White, and the Black identity category as implicitly masculine. To our knowledge, this is the first method that can probe word embeddings for these intersectional biases directly. Aside from addressing the issues discussed in Section 5, future work should investigate whether this method can be used to detect other sorts of intersectional biases, particularly those facing the LGBTQ community.

## References

- Samira Abnar, Lisa Beinborn, Rochelle Choenni, and Willem Zuidema. 2019. Blackbox meets blackbox: Representational similarity & stability analysis of neural language models and brains. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 191–203, Florence, Italy, August. Association for Computational Linguistics.
- Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013.
- Diane Bouchacourt and Marco Baroni. 2018. How agents see things: On visual representations in an emergent language game. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 981–985, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Lisa Bowleg. 2012. The problem with the phrase women and minorities: intersectionality—an important theoretical framework for public health. *American journal of public health*, 102(7):1267–1273.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- United States Census Bureau. 2020. 2020 census questions: Race. *2020Census.gov*.
- Grzegorz Chrupała and Afra Alishahi. 2019. Correlating neural and symbolic representations of language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2952–2962.
- Grzegorz Chrupała. 2019. Symbolic inductive bias for visually grounded learning of spoken language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6452–6462, Florence, Italy, July. Association for Computational Linguistics.
- Kimberlé Crenshaw. 1989. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *The University of Chicago Legal Forum*, page 139.
- Kimberlé Crenshaw. 1990. Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43:1241.
- 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.
- Roland G Fryer Jr and Steven D Levitt. 2004. The causes and consequences of distinctively black names. *The Quarterly Journal of Economics*, 119(3):767–805.
- S Michael Gaddis. 2017. How black are Lakisha and Jamal? Racial perceptions from names used in correspondence audit studies. *Sociological Science*, 4:469–489.
- 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.
- Anthony G Greenwald, Debbie E McGhee, and Jordan LK Schwartz. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464.
- Wei Guo and Aylin Caliskan. 2020. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. *arXiv preprint arXiv:2006.03955*.
- Aurélie Herbelot, Eva Von Redecker, and Johanna Müller. 2012. Distributional techniques for philosophical enquiry. In *Proceedings of the 6th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 45–54.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. 2008. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- 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.

- Michael A. Lepori and R. Thomas McCoy. 2020. Picking BERT’s brain: Probing for linguistic dependencies in contextualized embeddings using representational similarity analysis. In *Proceedings of 28th International Conference on Computational Linguistics*.
- Chandler May, Alex Wang, Shikha Bordia, Samuel 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.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2019. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*.
- Brian A Nosek, Mahzarin R Banaji, and Anthony G Greenwald. 2002a. Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6(1):101.
- Brian A Nosek, Mahzarin R Banaji, and Anthony G Greenwald. 2002b. Math= male, me= female, therefore math $\neq$  me. *Journal of personality and social psychology*, 83(1):44.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Rachel Rudinger, Chandler May, and Benjamin Van Durme. 2017. Social bias in elicited natural language inferences. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 74–79.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. Gender bias in coreference resolution. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14.
- 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.
- Latanya Sweeney. 2013. Discrimination in online ad delivery. *Queue*, 11(3):10–29.
- Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. In *Advances in Neural Information Processing Systems*, pages 13209–13220.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender bias in contextualized word embeddings. 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 629–634, Minneapolis, Minnesota, June. Association for Computational Linguistics.

## A Data

### A.1 Proper Names

We use the proper name data found in Sweeney (2013). That work largely compiles these proper names from two previous studies, which rely on correlations between race and child names in birth records from the years 1974-1979 in Massachusetts (Bertrand and Mullainathan, 2004), and in birth records from 1961-2000 in California (Fryer Jr and Levitt, 2004). We omit names that are polysemous. This occurs for four of the Black female names: *Diamond*, *Precious*, *Kenya*, and *Ebony*. One name (*Trevon*) did not occur in the Wikipedia 2014 + Gigaword 5 corpus, and so it was also excluded from our analysis. We randomly omit names from the Black male, White female, and White male proper name sets so that all identity categories are represented by an equal number (13) of names. The list of names from Sweeney (2013) is found below. Italicized names were omitted.

**Black Female:** Aisha, Keisha, Latonya, Lakisha, Latoya, Tamika, Imani, Shanice, Aaliyah, Nia, Latanya, Latisha, Deja, *Diamond*, *Precious*, *Kenya*, *Ebony*

**Black Male:** Darnell, Hakim, Jermaine, Kareem, Jamal, Leroy, Rasheed, DeShawn, DeAndre, Marquis, Terrell, Malik, Tyrone, *Tremayne*, *Trevon*

**White Female:** Allison, Anne, Carrie, Emily, Jill, Laurie, Kristen, Meredith, Molly, Amy, Claire, Madeline, Emma, *Katie*

**White Male:** Brad, Brendan, Geoffrey, Greg, Brett, Jay, Matthew, Jake, Connor, Tanner, Wyatt, Cody, Dustin, *Neil*

### A.2 Female Concept Words

We combine the datasets of words associated with the female identity category from Nosek et al. (2002a) and Nosek et al. (2002b), which are two often-cited studies that leverage the implicit association test (Greenwald et al., 1998). The list is as follows:

**Female Concept Words:** female, woman, girl, sister, she, her, hers, daughter, aunt, mother, grandmother

### A.3 Black Concept Words

Our Black concept words are derived from the definitions given by the United States Census Bureau (Census Bureau, 2020). The list is as follows:

**Black Concept Words:** Africa, Black, Jamaica, Haiti, Nigeria, Ethiopia, Somalia, Ghana, Barbados, Kenya, Liberia, Bahamas

### A.4 Multi-Word Group Items

Our group items leverage the semantically-bleached sentence templates of May et al. (2019). We list some examples of group items for all identity groups studies.

**Black Female** This is a Black female.; The Black woman is here.; They are Black girls.

**White Female** This is a White female.; The White woman is here.; They are White girls.

**Black Male** This is a Black male.; The Black man is here.; They are Black boys.

**White Male** This is a White male.; The White man is here.; They are White boys.

### A.5 Multi-Word Female Concept Items

Our multi-word female concept items leverage the semantically-bleached sentence templates of May et al. (2019). We list some examples of Female concept items.

**Female Concept Items** Females are people.; The girl is there.; Here is a woman.

## **A.6 Multi-Word Black Concept Items**

Our multi-word Black concept items leverage the semantically-bleached sentence templates of May et al. (2019). We list some examples of Black concept items.

**Black Concept Items** Here is a Black person.; These are Black people.; Black people are people.