



Citation for published version:

Caliskan, A, Bryson, JJ & Narayanan, A 2017, 'Semantics derived automatically from language corpora contain human-like biases', *Science*, vol. 356, no. 6334, pp. 183-186. <https://doi.org/10.1126/science.aal4230>

DOI:

[10.1126/science.aal4230](https://doi.org/10.1126/science.aal4230)

Publication date:

2017

Document Version

Peer reviewed version

[Link to publication](#)

This is the author's version of the work. It is posted here by permission of the AAAS for personal use, not for redistribution. The definitive version was published in *Science* on Vol. 356, Issue 6334, 14 April 2017, DOI: 10.1126/science.aal4230

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Semantics derived automatically from language corpora contain human-like biases

Aylin Caliskan¹, Joanna J. Bryson^{1,2}, Arvind Narayanan¹

¹Princeton University ²University of Bath

Machine learning is a means to derive artificial intelligence by discovering patterns in existing data. Here we show that applying machine learning to ordinary human language results in human-like semantic biases. We replicate a spectrum of known biases, as measured by the Implicit Association Test, using a widely used, purely statistical machine-learning model trained on a standard corpus of text from the Web. Our results indicate that text corpora contain recoverable and accurate imprints of our historic biases, whether morally neutral as towards insects or flowers, problematic as towards race or gender, or even simply veridical, reflecting the *status quo* distribution of gender with respect to careers or first names. Our methods hold promise for identifying and addressing sources of bias in culture, including technology.

Introduction

We show that standard machine learning can acquire stereotyped biases from textual data that reflect everyday human culture. The general idea that text corpora capture semantics including cultural stereotypes and empirical associations has long been known in corpus linguistics (1, 2), but our findings add to this knowledge in three ways. First, we use *word embeddings* (3), a

powerful tool to extract associations captured in text corpora; this method substantially amplifies the signal found in raw statistics. Second, our replication of documented human biases may yield tools and insights for studying prejudicial attitudes and behavior in humans. Third, since we perform our experiments on off-the-shelf machine learning components (primarily the GloVe word embedding), we show that cultural stereotypes propagate to Artificial Intelligence (AI) technologies in widespread use.

Before presenting our results, we discuss key terms and describe the tools we use. Terminology varies by discipline; these definitions are intended for clarity of the present article. In AI and machine learning, *bias* refers generally to prior information, a necessary prerequisite for intelligent action (4). Yet bias can be problematic where such information is derived from aspects of human culture known to lead to harmful behavior. Here we will call such biases ‘stereotyped’, and actions taken on their basis ‘prejudiced’.

We use the Implicit Association Test (IAT) as our primary source of documented human biases (5). The IAT demonstrates enormous differences in response times when subjects are asked to pair two concepts they find similar, in contrast to two concepts they find different. We develop a statistical test analogous to the IAT and apply it to a widely-used semantic representation of words in AI, termed *word embeddings*. Word embeddings represent each word as a vector in a vector space of around 300 dimensions, based on the textual context in which the word is found. We use the distance between a pair of vectors (more precisely, their cosine similarity score, a measure of correlation) as analogous to reaction time in the IAT.

Most closely related to this paper is concurrent work by Bolukbasi et al. (6), who propose a method to “debias” word embeddings. Our work is complementary, as we focus instead on rigorously demonstrating human-like biases in word embeddings. Further, our methods do not require an algebraic formulation of bias, which may not be possible for all types of bias. Additionally, we study the relationship between stereotyped associations and empirical data

concerning contemporary society.

Results

Using the measure of semantic association described above, we have been able to replicate every stereotype that we tested. We selected IATs that studied general societal attitudes, rather than those of subpopulations, and for which lists of target and attribute words (rather than images) were available. The results are summarized in Table 1.

Greenwald et al. introduce and validate the IAT by studying biases that they consider near-universal in humans and about which there is no social concern (5, p. 1469). We begin by replicating these inoffensive results for the same purposes. Specifically, they demonstrate that flowers are significantly more pleasant than insects, based on the reaction latencies of four pairings (flowers + pleasant, insects + unpleasant, flowers + unpleasant, insects + pleasant). Greenwald et al. measure effect size in terms of Cohen’s d , which is the difference between two means of log-transformed latencies in milliseconds, divided by the standard deviation. Conventional small, medium, and large values of d are 0.2, 0.5, and 0.8, respectively. With 32 participants, the IAT comparing flowers and insects results in an effect size of 1.35 ($p < 10^{-8}$). Applying our method, we observe the same expected association with an effect size of 1.50 ($p < 10^{-7}$). Similarly, we replicate Greenwald et al.’s finding [p. 1469] that musical instruments are significantly more pleasant than weapons (See Table 1).

Notice that the word embeddings “know” these properties of flowers, insects, musical instruments, and weapons with no direct experience of the world, and no representation of semantics other than the implicit metrics of words’ co-occurrence statistics with other nearby words.

We now use the same technique to demonstrate that machine learning absorbs stereotyped biases as easily as any other. Greenwald et al. (5, p. 1475) find extreme impacts of race as indicated simply by name. A bundle of names associated with being *European American* was

Target words	Attrib. words	Original Finding				Our Finding			
		Ref	N	d	p	N _T	N _A	d	p
Flowers vs insects	Pleasant vs unpleasant	(5)	32	1.35	10 ⁻⁸	25×2	25×2	1.50	10 ⁻⁷
Instruments vs weapons	Pleasant vs unpleasant	(5)	32	1.66	10 ⁻¹⁰	25×2	25×2	1.53	10 ⁻⁷
Eur.-American vs Afr.-American names	Pleasant vs unpleasant	(5)	26	1.17	10 ⁻⁵	32×2	25×2	1.41	10 ⁻⁸
Eur.-American vs Afr.-American names	Pleasant vs unpleasant from (5)	(7)	Not applicable			16×2	25×2	1.50	10 ⁻⁴
Eur.-American vs Afr.-American names	Pleasant vs unpleasant from (9)	(7)	Not applicable			16×2	8×2	1.28	10 ⁻³
Male vs female names	Career vs family	(9)	39k	0.72	< 10 ⁻²	8×2	8×2	1.81	10 ⁻³
Math vs arts	Male vs female terms	(9)	28k	0.82	< 10 ⁻²	8×2	8×2	1.06	.018
Science vs arts	Male vs female terms	(10)	91	1.47	10 ⁻²⁴	8×2	8×2	1.24	10 ⁻²
Mental vs physical disease	Temporary vs permanent	(23)	135	1.01	10 ⁻³	6×2	7×2	1.38	10 ⁻²
Young vs old people’s names	Pleasant vs unpleasant	(9)	43k	1.42	< 10 ⁻²	8×2	8×2	1.21	10 ⁻²

Table 1: Summary of Word Embedding Association Tests. We replicate 8 well-known IAT findings using word embeddings (rows 1–3 and 6–10); we also help explain prejudiced human behavior concerning hiring in the same way (rows 4 and 5). Each result compares two sets of words from target *concepts* about which we are attempting to learn with two sets of *attribute* words. In each case the first target is found compatible with the first attribute, and the second target with the second attribute. Throughout, we use word lists from the studies we seek to replicate. N : number of subjects. N_T : number of target words. N_A : number of attribute words. We report the effect sizes (d) and p -values (p , rounded up) to emphasize that the statistical and substantive significance of both sets of results is uniformly high; we do not imply that our numbers are directly comparable to those of human studies. For the online IATs (rows 6, 7, and 10), p -values were not reported, but are known to be below the significance threshold of 10⁻². Rows 1–8 are discussed in the text; for completeness, this table also includes the two other IATs for which we were able to find suitable word lists (rows 9 and 10).

found to be significantly more easily associated with pleasant than unpleasant terms, compared to a bundle of *African American* names.

In replicating this result, we were forced to slightly alter the stimuli because some of the original African American names did not occur in the corpus with sufficient frequency to be included. We therefore also deleted the same number of European American names, chosen at random, to balance the number of elements in the sets of two concepts. Omissions and deletions are indicated in our list of keywords (see Supplement).

In another widely-publicized study, Bertrand and Mullainathan (7) sent nearly 5,000 identical résumés in response to 1,300 job advertisements, varying only the names of the candidates. They found that European American candidates were 50% more likely to be offered an opportunity to be interviewed. In follow-up work, they argue that implicit biases help account for these effects (8).

We provide additional evidence for this hypothesis using word embeddings. We test the names in their study for pleasantness associations. As before, we had to delete some low-frequency names. We confirm the association using two different sets of ‘pleasant/unpleasant’ stimuli: those from the original IAT paper, and also a shorter, revised set published later (9).

Turning to gender biases, we replicate a finding that *female names* are more associated with family than career words, compared to *male names* (9). This IAT was conducted online, and thus has a vastly larger subject pool, but far fewer keywords. We replicate the IAT results even with these reduced keyword sets. We also replicate an online IAT finding that *female words* (“woman”, “girl”, ...) are more associated than *male words* with the arts than mathematics (9, p. 105). Finally, we replicate a laboratory study showing that female words are more associated with the arts than the sciences (10, p. 51).

Comparison to Real-World Data. It has been suggested that implicit gender-occupation biases are linked to gender gaps in occupational participation; however the relationship between

these is complex and may be mutually reinforcing (11).




Figure 1: Occupation-gender association. Pearson's correlation coefficient $\rho = 0.90$ with p -value $< 10^{-18}$.




Figure 2: Name-gender association. Pearson's correlation coefficient $\rho = 0.84$ with p -value $< 10^{-13}$.

To better understand the relationship, we examine the correlation between the gender association of occupation words and labor-force participation data. The x-axis of Figure 1 is derived from 2015 data released by the U.S. Bureau of Labor Statistics (<http://www.bls.gov/cps/cpsaat11.htm>), which provides information about occupational categories and the percentage of women that have certain occupations under these categories. By applying a second method that we developed, *Word Embedding Factual Association Test*, (WEFAT), we find that GloVe word embeddings correlate strongly with the percentage of women in 50 occupations in the USA in 2015.

Similarly, we looked at the veridical association of gender to androgynous names, that is, names used by either gender. In this case, the most recent information we were able to find was the 1990 census name and gender statistics. Perhaps because of the age of our name data, our correlation was weaker than for the 2015 occupation statistics, but still strikingly significant. In Fig 2, the x-axis is derived from the 1990 U.S. census data (<http://www.census.gov/main/www/cen1990.html>), and the y-axis is as before.

Methods

A word embedding is a representation of words as points in a vector space (12). For all results in this paper we use the state-of-the-art GloVe word embedding method, in which, at a high level, the similarity between a pair of vectors is related to the probability that the words co-occur with other words similarly to each other in text (13). Word embedding algorithms such as GloVe exploit dimensionality reduction to substantially amplify the signal found in simple co-occurrence probabilities. In pilot-work experiments along the lines of those presented here (on free associations rather than implicit associations) raw co-occurrence probabilities were shown to lead to much weaker results (14, 15).

Rather than train the embedding ourselves, we use pre-trained GloVe embeddings distributed by its authors. This ensures impartiality, simplifies reproducing our results, and allows us to replicate the effects that may be found in real applications of machine learning. We use the largest of the four corpora provided—the “Common Crawl” corpus obtained from a large-scale crawl of the web, containing 840 billion tokens (roughly, words). Tokens in this corpus are case-sensitive, resulting in 2.2 million different ones. Each word corresponds to a 300-dimensional vector derived from counts of other words that co-occur with it in a 10-word window.

In the Supplement we also present substantially similar results using an alternative corpus and word embedding.

Word Embedding Association Test (WEAT). Borrowing terminology from the IAT literature, consider two sets of target words (e.g., programmer, engineer, scientist, ... and nurse, teacher, librarian, ...) and two sets of *attribute* words (e.g., man, male, ... and woman, female ...). The null hypothesis is that there is no difference between the two sets of target words in terms of their relative similarity to the two sets of attribute words. The permutation test measures the (un)likelihood of the null hypothesis by computing the probability that a random

permutation of the attribute words would produce the observed (or greater) difference in sample means.

In formal terms, let X and Y be two sets of target words of equal size, and A, B the two sets of attribute words. Let $\cos(\vec{a}, \vec{b})$ denote the cosine of the angle between the vectors \vec{a} and \vec{b} .

The test statistic is

$$s(X, Y, A, B) = \sum_{x \in X} s(x, A, B) - \sum_{y \in Y} s(y, A, B) \quad \text{where}$$

$$s(w, A, B) = \text{mean}_{a \in A} \cos(\vec{w}, \vec{a}) - \text{mean}_{b \in B} \cos(\vec{w}, \vec{b})$$

In other words, $s(w, A, B)$ measures the association of w with the attribute, and $s(X, Y, A, B)$ measures the differential association of the two sets of target words with the attribute.

Let $\{(X_i, Y_i)\}_i$ denote all the partitions of $X \cup Y$ into two sets of equal size. The one-sided p -value of the permutation test is

$$\Pr_i[s(X_i, Y_i, A, B) > s(X, Y, A, B)]$$

The effect size is

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

This is a normalized measure of how separated the two distributions (of associations between the target and attribute) are. We re-iterate that these p -values and effect sizes don't have the same interpretation as the IAT, as the "subjects" in our experiments are words, not people.

Word Embedding Factual Association Test (WEFAT). This test helps us examine how word embeddings capture empirical information about the world embedded in text corpora. Consider a set of target concepts, such as occupations, and a real-valued, factual property of the world associated with each concept, such as the percentage of workers in the occupation who are women. We'd like to test if the vectors corresponding to the concepts embed knowledge of the property, that is, if there is an algorithm that can extract or predict the property given the

vector. In principle we could use any algorithm, but in this work we test the association of the target concept with some set of attribute words, analogous to WEAT.

Formally, consider a single set of target words W and two sets of attribute words A, B . There is a property p_w associated with each word $w \in W$.

- The statistic associated with each word vector is a normalized association score of the word with the attribute:

$$s(w, A, B) = \frac{\text{mean}_{a \in A} \cos(\vec{w}, \vec{a}) - \text{mean}_{b \in B} \cos(\vec{w}, \vec{b})}{\text{std-dev}_{x \in A \cup B} \cos(\vec{w}, \vec{x})}$$

- The null hypothesis is that there is no association between $s(w, A, B)$ and p_w . We test the null hypothesis using a linear regression analysis to predict the latter from the former.

Discussion

We elaborate on further implications of our results. In psychology, our results add to the credence of the IAT by replicating its results in such a different setting. Further, our methods may yield an efficient way to explore previously unknown implicit associations. Researchers who conjecture an implicit association might first test them using WEAT on a suitable corpus before testing human subjects. Similarly, our methods could be used to quickly find differences in bias between demographic groups, given large corpora authored by members of the respective groups. If substantiated through testing and replication, WEAT may also give us access to implicit associations of groups not available for testing, such as historic populations.

We have demonstrated that word embeddings encode not only stereotyped biases but also other knowledge, such as the visceral pleasantness of flowers or the gender distribution of occupations. These results lend support to the *distributional hypothesis* in linguistics, namely that the statistical contexts of words capture much of what we mean by meaning (16). Our findings

are also sure to contribute to the debate concerning the Sapir-Whorf hypothesis (17), since our work suggests that behavior can be driven by cultural history embedded in a term's historic use. Such histories can evidently vary between languages.

We stress that we replicated every association documented via the IAT that we tested. The number, variety, and substantive significance of our results raise the possibility that all implicit human biases are reflected in the statistical properties of language. Further research is needed to test this hypothesis and to compare language to other modalities, especially the visual, to see if they have similarly strong explanatory power.

Our results also suggest a null hypothesis for explaining origins of prejudicial behavior in humans, namely, the implicit transmission of ingroup/outgroup identity information through language. That is, before providing an explicit or institutional explanation for why individuals make prejudiced decisions, one must show that it was not a simple outcome of unthinking reproduction of statistical regularities absorbed with language. Similarly, before positing complex models for how stereotyped attitudes perpetuate from one generation to the next or from one group to another, we must check whether simply learning language is sufficient to explain (some of) the observed transmission of prejudice.

Our work has implications for AI and machine learning because of the concern that these technologies may perpetuate cultural stereotypes (18). Our findings suggest that if we build an intelligent system that learns enough about the properties of language to be able to understand and produce it, in the process it will also acquire historic cultural associations, some of which can be objectionable. Already, popular online translation systems incorporate some of the biases we study; see Supplement. Further concerns may arise as AI is given agency in our society. If machine learning technologies used for, say, résumé screening were to imbibe cultural stereotypes, it may result in prejudiced outcomes. We recommend addressing this through the explicit characterization of acceptable behavior. One such approach is seen in the nascent

field of fairness in machine learning, which specifies and enforces mathematical formulations of non-discrimination in decision-making (19, 20). Another approach can be found in modular AI architectures, such as cognitive systems, in which implicit learning of statistical regularities can be compartmentalized and augmented with explicit instruction of rules of appropriate conduct (21, 22). Certainly, caution must be used in incorporating modules constructed via unsupervised machine learning into decision-making systems.

References and Notes

1. M. Stubbs, *Text and corpus analysis: Computer-assisted studies of language and culture* (Blackwell Oxford, 1996).
2. J. A. Bullinaria, J. P. Levy, Extracting semantic representations from word co-occurrence statistics: A computational study, *Behavior Research Methods* **39**, 510 (2007).
3. T. Mikolov, J. Dean, Distributed representations of words and phrases and their compositionality, *Advances in neural information processing systems* (2013).
4. C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, London, 2006).
5. A. G. Greenwald, D. E. McGhee, J. L. Schwartz, Measuring individual differences in implicit cognition: the implicit association test., *Journal of personality and social psychology* **74**, 1464 (1998).
6. T. Bolukbasi, K.-W. Chang, J. Y. Zou, V. Saligrama, A. T. Kalai, Man is to computer programmer as woman is to homemaker? Debiasing word embeddings, *Advances in Neural Information Processing Systems* (2016), pp. 4349–4357.
7. M. Bertrand, S. Mullainathan, Are Emily and Greg more employable than Lakisha and Jamal? a field experiment on labor market discrimination, *The American Economic Review* **94**, 991 (2004).
8. M. Bertrand, D. Chugh, S. Mullainathan, Implicit discrimination, *American Economic Review* pp. 94–98 (2005).
9. B. A. Nosek, M. Banaji, A. G. Greenwald, Harvesting implicit group attitudes and beliefs from a demonstration web site., *Group Dynamics: Theory, Research, and Practice* **6**, 101 (2002).

10. B. A. Nosek, M. R. Banaji, A. G. Greenwald, Math=male, me=female, therefore math≠me., *Journal of Personality and Social Psychology* **83**, 44 (2002).
11. B. A. Nosek, *et al.*, National differences in gender–science stereotypes predict national sex differences in science and math achievement, *Proceedings of the National Academy of Sciences* **106**, 10593 (2009).
12. P. D. Turney, P. Pantel, From frequency to meaning: Vector space models of semantics, *Journal of Artificial Intelligence Research* **37**, 141 (2010).
13. J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation., *EMNLP* (2014), vol. 14, pp. 1532–43.
14. T. MacFarlane, Extracting semantics from the Enron corpus, *University of Bath, Department of Computer Science Technical Report Series; CSBU-2013-08* <http://opus.bath.ac.uk/37916/> (2013).
15. W. Lowe, S. McDonald, The direct route: Mediated priming in semantic space, *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society* (LEA, 2000), pp. 806–811.
16. M. Sahlgren, The distributional hypothesis, *Italian Journal of Linguistics* **20**, 33 (2008).
17. G. Lupyan, The centrality of language in human cognition, *Language Learning* **66**, 516 (2016).
18. S. Barocas, A. D. Selbst, Big data’s disparate impact, *California Law Review* **104** (2014).
19. C. Dwork, M. Hardt, T. Pitassi, O. Reingold, R. Zemel, Fairness through awareness, *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (ACM, 2012), pp. 214–226.

20. M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, S. Venkatasubramanian, Certifying and removing disparate impact, *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, 2015), pp. 259–268.
21. K. R. Thórisson, Integrated A.I. systems, *Minds and Machines* **17**, 11 (2007).
22. M. Hanheide, *et al.*, Robot task planning and explanation in open and uncertain worlds, *Artificial Intelligence* (2015).
23. L. L. Monteith, J. W. Pettit, Implicit and explicit stigmatizing attitudes and stereotypes about depression, *Journal of Social and Clinical Psychology* **30**, 484 (2011).

Acknowledgements. We are grateful to the following people: Will Lowe for substantial assistance in the design of our significance tests, Tim Macfarlane for his pilot research as a part of his undergraduate dissertation, Solon Barocas, Miles Brundage, Kate Crawford, Calvin Lai, and Matthew Salganik for extremely useful comments on a draft of this paper. We have archived the code and data on Harvard Dataverse (doi:10.7910/DVN/DX4VWP).

Supplementary Material

List of supplementary content:

- Materials and methods, providing additional details of our method.
- Text, figures, and legends, including replication of our results with an alternate corpus and algorithm, results on stereotypes reflected in statistical machine translation, and the list of stimuli used in our analyses.
- Tables and legends, containing a summary of replication results with an alternate corpus and algorithm.

Materials and Methods

Cosine similarity. Given two vectors $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$ and $\mathbf{y} = \langle y_1, y_2, \dots, y_n \rangle$, their cosine similarity can be calculated as:

$$\cos(\mathbf{x}, \mathbf{y}) = \frac{\sum_{i=1}^n x_i \cdot y_i}{\sqrt{\sum_{i=1}^n x_i^2} \sqrt{\sum_{i=1}^n y_i^2}}$$

In other words, it is the dot product of the vectors after they have been normalized to unit length.

Applying the Word Embedding Factual Association Test (WEFAT). Now we discuss in more detail how we apply WEFAT in two cases. The first is to test if occupation word vectors embed knowledge of the gender composition of the occupation in the real world. We use data released by the Bureau of Labor Statistics in which occupations are categorized hierarchically, and for each occupation the number of workers and percentage of women are given (some data is missing). The chief difficulty is that many occupation names are multi-word terms whereas the pre-trained word vectors that we use represent single words. Our strategy is to convert a

multi-word term into a single word that represents a superset of the category (e.g., *chemical engineer* → *engineer*), and filter out occupations where this is not possible. The resulting words are listed in the following section.

Our second application of WEFAT is to test if androgynous names embed knowledge of how often the name is given to boys versus girls. We picked the most popular names in each 10% window of gender frequency based on 1990 U.S. Census data. Here again there is a difficulty: some names are also regular English words (e.g., *Will*). State-of-the-art word embeddings are not yet sophisticated enough to handle words with multiple senses or meanings; all usages are lumped into a single vector. To handle this, we algorithmically determine how “name-like” each vector is (by computing the distance of each vector to the centroid of all the name vectors), and eliminate the 20% of vectors that are least name-like.

Caveats about comparing WEAT to the IAT. In WEAT, much like the IAT, we do not compare two words. Many if not most words have multiple meanings, which makes pairwise measurements “noisy”. To control for this, we use small baskets of terms to represent a concept. In every case we use word baskets from previous psychological studies, typically from the same study we are replicating. We should note that distances / similarities of word embeddings lack any intuitive interpretation. But this poses no problem for us: our results and their import do not depend on attaching meaning to these distances.

While the IAT applies to individual human subjects, the embeddings of interest to us are derived from the *aggregate* writings of humans on the web. These corpora are generated in an uncontrolled fashion and are not representative of any one population. The IAT has been used to draw conclusions about populations by averaging individual results over samples. Our tests of word embeddings are loosely analogous to such population-level IATs.

Nevertheless, this difference precludes a direct numerical comparison between human biases measured by the IAT and biases in corpora measured by our methods. With word embed-

dings, there is no notion of test subjects. Roughly, it is as if we are able to measure the mean of the association strength over all the “subjects” who collectively created the corpus. But we have no way to observe variation between subjects or between trials. We do report p -values and effect sizes resulting from the use of multiple *words* in each category, but the meaning of these numbers is entirely different from those reported in IATs.

Text, figures, and legends

Replicating our results with other corpora and algorithms. We repeated all the *WEAT* and *WEFAT* analyses presented above using a different pre-trained embedding: word2vec on a Google News corpus (3). The embedding contains 3 million word vectors, and the corpus contains about 100 billion tokens, about an order of magnitude smaller than the Common Crawl corpus. Therefore the less common terms (especially names) in our lists occur infrequently in this corpus. This makes replication harder, as the co-occurrence statistics are “noisier”. Yet in all WEATs except one, we observed statistically significant effects ($p < .05$) and large effect sizes. The lone exception is the pleasantness association of young vs. old people’s names, a test which has a small number of target concepts and relatively low keyword frequencies. Table S1 summarizes the results.

Further, we found that the gender association strength of occupation words is highly correlated between the GloVe embedding and the word2vec embedding (Pearson $\rho = 0.88$; Spearman $\rho = 0.86$). In concurrent work, Bolukbasi et al. (6) compared the same two embeddings, using a different measure of the gender bias of occupation words, also finding a high correlation (Spearman $\rho = 0.81$).

Stereotypes reflected in statistical machine translation. One application where we can observe cultural stereotypes reflected is Statistical machine translation (SMT), a common natural language processing task. Translations to English from many gender-neutral languages

such as Turkish lead to gender-stereotyped sentences. For example, Google Translate converts these Turkish sentences with gender-neutral pronouns: “O bir doktor. O bir hemşire.” to these English sentences: “He is a doctor. She is a nurse.” We see the same behavior for Finnish, Estonian, Hungarian, and Persian in place of Turkish. Similarly, translating the above two Turkish sentences into several of the most commonly spoken languages (Spanish, English, Portuguese, Russian, German, and French) results in gender-stereotyped pronouns in every case.

List of stimuli. Here we list the stimuli used in our WEAT and WEFAT tests. The WEAT tests are listed in the same order as Table 1.

WEAT 1: We use the flower and insect target words along with pleasant and unpleasant attributes found in (5).

- Flowers: aster, clover, hyacinth, marigold, poppy, azalea, crocus, iris, orchid, rose, bluebell, daffodil, lilac, pansy, tulip, buttercup, daisy, lily, peony, violet, carnation, gladiola, magnolia, petunia, zinnia.
- Insects: ant, caterpillar, flea, locust, spider, bedbug, centipede, fly, maggot, tarantula, bee, cockroach, gnat, mosquito, termite, beetle, cricket, hornet, moth, wasp, blackfly, dragonfly, horsefly, roach, weevil.
- Pleasant: caress, freedom, health, love, peace, cheer, friend, heaven, loyal, pleasure, diamond, gentle, honest, lucky, rainbow, diploma, gift, honor, miracle, sunrise, family, happy, laughter, paradise, vacation.
- Unpleasant: abuse, crash, filth, murder, sickness, accident, death, grief, poison, stink, assault, disaster, hatred, pollute, tragedy, divorce, jail, poverty, ugly, cancer, kill, rotten, vomit, agony, prison.

WEAT 2: We use the musical instruments and weapons target words along with pleasant and unpleasant attributes found in (5).

- Instruments: bagpipe, cello, guitar, lute, trombone, banjo, clarinet, harmonica, mandolin, trumpet, bassoon, drum, harp, oboe, tuba, bell, fiddle, harpsichord, piano, viola, bongo, flute, horn, saxophone, violin.
- Weapons: arrow, club, gun, missile, spear, axe, dagger, harpoon, pistol, sword, blade, dynamite, hatchet, rifle, tank, bomb, firearm, knife, shotgun, teargas, cannon, grenade, mace, slingshot, whip.
- Pleasant: As per previous experiment with insects and flowers.
- Unpleasant: As per previous experiment with insects and flowers.

WEAT 3: We use the European American and African American names along with pleasant and unpleasant attributes found in (5). Names that are marked with italics are excluded from our replication. In the case of African American names this was due to being too infrequent to occur in GloVe's Common Crawl corpus; in the case of European American names an equal number were deleted, chosen at random.

- European American names: Adam, *Chip*, Harry, Josh, Roger, Alan, Frank, *Ian*, Justin, Ryan, Andrew, *Fred*, Jack, Matthew, Stephen, Brad, Greg, *Jed*, Paul, *Todd*, *Brandon*, *Hank*, Jonathan, Peter, *Wilbur*, Amanda, Courtney, Heather, Melanie, *Sara*, *Amber*, *Crystal*, *Katie*, *Meredith*, *Shannon*, Betsy, *Donna*, Kristin, Nancy, Stephanie, *Bobbie-Sue*, Ellen, Lauren, *Peggy*, *Sue-Ellen*, Colleen, Emily, Megan, Rachel, *Wendy* (deleted names in italics).
- African American names: Alonzo, Jamel, *Lerone*, *Percell*, Theo, Alphonse, Jerome, Leroy, *Rasaan*, Torrance, Darnell, Lamar, Lionel, *Rashaun*, Tyree, Deion, Lamont, Malik,

Terrence, Tyrone, *Everol*, Lavon, Marcellus, *Terryl*, Wardell, *Aiesha*, *Lashelle*, Nichelle, Shereen, *Temeka*, Ebony, Latisha, Shaniqua, *Tameisha*, *Teretha*, Jasmine, *Latonya*, *Shanise*, Tanisha, Tia, Lakisha, Latoya, *Sharise*, *Tashika*, Yolanda, *Lashandra*, Malika, *Shavonn*, *Tawanda*, Yvette (deleted names in italics).

- Pleasant: caress, freedom, health, love, peace, cheer, friend, heaven, loyal, pleasure, diamond, gentle, honest, lucky, rainbow, diploma, gift, honor, miracle, sunrise, family, happy, laughter, paradise, vacation.
- Unpleasant: abuse, crash, filth, murder, sickness, accident, death, grief, poison, stink, assault, disaster, hatred, pollute, tragedy, bomb, divorce, jail, poverty, ugly, cancer, evil, kill, rotten, vomit.

WEAT 4: We use the European American and African American names from (7), along with pleasant and unpleasant attributes found in (5).

- European American names: Brad, Brendan, Geoffrey, Greg, Brett, *Jay*, Matthew, Neil, Todd, Allison, Anne, Carrie, Emily, Jill, Laurie, *Kristen*, Meredith, Sarah (names in italics deleted in GloVe experiments).
- African American names: Darnell, Hakim, Jermaine, Kareem, Jamal, Leroy, Rasheed, *Tremayne*, Tyrone, Aisha, Ebony, Keisha, Kenya, *Latonya*, Lakisha, Latoya, Tamika, Tanisha (names in italics deleted in GloVe experiments).
- Pleasant: caress, freedom, health, love, peace, cheer, friend, heaven, loyal, pleasure, diamond, gentle, honest, lucky, rainbow, diploma, gift, honor, miracle, sunrise, family, happy, laughter, paradise, vacation.

- Unpleasant: abuse, crash, filth, murder, sickness, accident, death, grief, poison, stink, assault, disaster, hatred, pollute, tragedy, bomb, divorce, jail, poverty, ugly, cancer, evil, kill, rotten, vomit.

WEAT 5: We use the European American and African American names from (7), along with pleasant and unpleasant attributes found in (9).

- European American names: Brad, Brendan, Geoffrey, Greg, Brett, *Jay*, Matthew, Neil, Todd, Allison, Anne, Carrie, Emily, Jill, Laurie, *Kristen*, Meredith, Sarah (names in italics deleted in GloVe experiments).
- African American names: Darnell, Hakim, Jermaine, Kareem, Jamal, Leroy, Rasheed, *Tremayne*, Tyrone, Aisha, Ebony, Keisha, Kenya, *Latonya*, Lakisha, Latoya, Tamika, Tanisha (names in italics deleted in GloVe experiments).
- Pleasant: joy, love, peace, wonderful, pleasure, friend, laughter, happy.
- Unpleasant: agony, terrible, horrible, nasty, evil, war, awful, failure.

WEAT 6: We use the male and female names along with career and family attributes found in (9).

- Male names: John, Paul, Mike, Kevin, Steve, Greg, Jeff, Bill.
- Female names: Amy, Joan, Lisa, Sarah, Diana, Kate, Ann, Donna.
- Career: executive, management, professional, corporation, salary, office, business, career.
- Family: home, parents, children, family, cousins, marriage, wedding, relatives.

WEAT 7: We use the math and arts target words along with male and female attributes found in (9).

- Math: math, algebra, geometry, calculus, equations, computation, numbers, addition.
- Arts: poetry, art, dance, literature, novel, symphony, drama, sculpture.
- Male terms: male, man, boy, brother, he, him, his, son.
- Female terms: female, woman, girl, sister, she, her, hers, daughter.

WEAT 8: We use the science and arts target words along with male and female attributes found in (10).

- Science: science, technology, physics, chemistry, Einstein, NASA, experiment, astronomy.
- Arts: poetry, art, Shakespeare, dance, literature, novel, symphony, drama.
- Male terms: brother, father, uncle, grandfather, son, he, his, him.
- Female terms: sister, mother, aunt, grandmother, daughter, she, hers, her.

WEAT 9: We use the mental and physical disease target words along with uncontrollability and controllability attributes found in (23).

- Mental disease: sad, hopeless, gloomy, tearful, miserable, depressed.
- Physical disease: sick, illness, influenza, disease, virus, cancer.
- Temporary: impermanent, unstable, variable, fleeting, *short-term*, brief, occasional (word2vec experiments used short instead of short-term).
- Permanent: stable, always, constant, persistent, chronic, prolonged, forever.

WEAT 10: We use young and old people's names as target words along with pleasant and unpleasant attributes found in (9).

- Young people's names: Tiffany, Michelle, Cindy, Kristy, Brad, Eric, Joey, Billy.
- Old people's names: Ethel, Bernice, Gertrude, Agnes, Cecil, Wilbert, Mortimer, Edgar.
- Pleasant: joy, love, peace, wonderful, pleasure, friend, laughter, happy.
- Unpleasant: agony, terrible, horrible, nasty, evil, war, awful, failure.

WEFAT 1 (occupations): We use the gender stimuli found in (9) along with the occupation attributes we derived from Bureau of Labor Statistics.

- **Careers :** technician, accountant, supervisor, engineer, worker, educator, clerk, counselor, inspector, mechanic, manager, therapist, administrator, salesperson, receptionist, librarian, advisor, pharmacist, janitor, psychologist, physician, carpenter, nurse, investigator, bartender, specialist, electrician, officer, pathologist, teacher, lawyer, planner, practitioner, plumber, instructor, surgeon, veterinarian, paramedic, examiner, chemist, machinist, appraiser, nutritionist, architect, hairdresser, baker, programmer, paralegal, hygienist, scientist.
- **Female attributes:** female, woman, girl, sister, she, her, hers, daughter.
- **Male attributes:** male, man, boy, brother, he, him, his, son.

WEFAT 2 (androgynous names): We use the gender stimuli found in (9) along with the most popular androgynous names from 1990's public census data as targets.

- **Names :** Kelly, Tracy, Jamie, Jackie, Jesse, Courtney, Lynn, Taylor, Leslie, Shannon, Stacey, Jessie, Shawn, Stacy, Casey, Bobby, Terry, Lee, Ashley, Eddie, Chris, Jody, Pat,

Carey, Willie, Morgan, Robbie, Joan, Alexis, Kris, Frankie, Bobbie, Dale, Robin, Billie, Adrian, Kim, Jaime, Jean, Francis, Marion, Dana, Rene, Johnnie, Jordan, Carmen, Ollie, Dominique, Jimmie, Shelby.

- **Female and Male attributes:** as per previous experiment on occupations.

Tables and legends

Target words	Attrib. words	Original Finding				Our Finding			
		Ref	N	d	p	N _T	N _A	d	p
Flowers vs insects	Pleasant vs unpleasant	(5)	32	1.35	10^{-8}	25×2	25×2	1.54	10^{-7}
Instruments vs weapons	Pleasant vs unpleasant	(5)	32	1.66	10^{-10}	25×2	25×2	1.63	10^{-8}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant	(5)	26	1.17	10^{-5}	32×2	25×2	0.58	10^{-2}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant	(7)	Not applicable			18×2	25×2	1.24	10^{-3}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant from (5)	(7)	Not applicable			18×2	8×2	0.72	10^{-2}
Male vs female names	Career vs family	(9)	39k	0.72	10^{-2}	8×2	8×2	1.89	10^{-4}
Math vs arts	Male vs female terms	(9)	28k	0.82	$< 10^{-2}$	8×2	8×2	0.97	.027
Science vs arts	Male vs female terms	(10)	91	1.47	10^{-24}	8×2	8×2	1.24	10^{-2}
Mental vs physical disease	Temporary vs permanent	(23)	135	1.01	10^{-3}	6×2	7×2	1.30	.012
Young vs old people's names	Pleasant vs unpleasant	(9)	43k	1.42	$< 10^{-2}$	8×2	8×2	-.08	0.57

Table S1: Summary of Word Embedding Association Tests using word2vec embeddings trained on the Google News corpus. The rows and columns are as in Table 1. For certain tests, the number of WEAT target words here is different than in Table 1, because in each case, we delete words not found in the corresponding word embedding.