



Sutton, A., Lansdall-Welfare, T., & Cristianini, N. (2018). Biased Embeddings from Wild Data: Measuring, Understanding and Removing. In *Advances in Intelligent Data Analysis XVII : 17th International Symposium, IDA 2018, 's-Hertogenbosch, The Netherlands, October 24–26, 2018, Proceedings* (pp. 328-339). (Lecture Notes in Computer Science; Vol. 11191). Springer, Cham. https://doi.org/10.1007/978-3-030-01768-2_27

Peer reviewed version

Link to published version (if available):
[10.1007/978-3-030-01768-2_27](https://doi.org/10.1007/978-3-030-01768-2_27)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Springer at https://link.springer.com/chapter/10.1007/978-3-030-01768-2_27 . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Biased Embeddings from Wild Data: Measuring, Understanding and Removing

Adam Sutton, Thomas Lansdall-Welfare, Nello Cristianini

Intelligent Systems Laboratory, University of Bristol, Bristol, BS8 1UB, UK
{adam.sutton, thomas.lansdall-welfare, nello.cristianini}@bris.ac.uk

Abstract. Many modern Artificial Intelligence (AI) systems make use of data embeddings, particularly in the domain of Natural Language Processing (NLP). These embeddings are learnt from data that has been gathered “from the wild” and have been found to contain unwanted biases. In this paper we make three contributions towards measuring, understanding and removing this problem. We present a rigorous way to measure some of these biases, based on the use of word lists created for social psychology applications; we observe how gender bias in occupations reflects actual gender bias in the same occupations in the real world; and finally we demonstrate how a simple projection can significantly reduce the effects of embedding bias. All this is part of an ongoing effort to understand how trust can be built into AI systems.

Keywords: Fairness in AI · Bias in Data · Artificial Intelligence · Natural Language Processing · Word Embeddings

1 Introduction

With the latest wave of learning models taking advantage of advances in deep learning [21,22,23], Artificial Intelligence (AI) systems are gaining widespread publicity, coupled with a drive from industry to incorporate intelligence into all manner of processes that handle our private and personal data, giving them a central position in our modern-day society.

This development has led to demand for fairer AI, where we wish to establish trust in the automated intelligent systems by ensuring that systems represent us fairly and transparently. However, there has been growing concern about potential biases in learning systems [1,6] which can be difficult to analyse or query for explanations of their predictions, leading to an increasing number of studies investigating the way black-box systems represent knowledge and make decisions [7,9,11,19,20]. Indeed, principled methods are now required that allow us to measure, understand and remove biases in our data in order for these systems to be truly accepted as a prominent part of our lives.

In the domain of text, many modern approaches often begin by embedding the input text data into an embedding space that is used as the first layer in a subsequent deep network [4,14]. These word embeddings have been shown to contain the same biases [3], due to the source data from which they are trained.

In effect, biases from the source data, such as in the differences in representation for men and women, that have been found in many different large-scale studies [5,10,12], carry through to the semantic relations in the word embeddings, which become baked into the learning systems that are built on top of them.

In this paper, we make three contributions towards addressing these concerns. First we propose a new version of the Word Embedding Association Tests (WEATs) studied in [3], designed to demonstrate and quantify bias in word embeddings, which puts them on a firm foundation by using the Linguistic Inquiry and Word Count (LIWC) lexica [17] to systematically *detect* and *measure* embedding biases.

With this improved experimental setting, we find that European-American names are viewed more positively than African-American names, male names are more associated with work while female names are more associated with family, and that the academic disciplines of science and maths are more associated with male terms than the arts, which are more associated with female terms. Using this new methodology, we then find that there is a gender bias in the way different occupations are represented by the embedding. Furthermore, we use the latest official employment statistics in the UK, and find that there is a correlation between the ratio of men and women working in different occupation roles and how those roles are associated with gender in the word embeddings. This suggests that biases in the embeddings reflect biases in the world.

Finally, we look at methods of *removing* gender bias from the word embeddings. Having established that there is a direction in the embedding space that correlates with gender, we use a simple orthogonal projection to remove that dimension from the embedding. After projecting the embeddings, we investigate the effect on bias in the embeddings by considering the changes in associations between the words, demonstrating that the associations in the modified embeddings now correlate less to UK employment statistics among other things.

2 Methodology

2.1 Word Embedding

A word embedding is a mapping of words into an n -dimensional vector space. Given a corpus of text, a word embedding can be created that will translate that corpus into a set of semantic vectors representing each word. Each word that appears in the corpus will be represented by an n -dimensional vector to indicate its position within the embedding.

This embedding has a set of features that can be used in natural language processing methods. The nearest neighbours of a word will be other words that have similar linguistic or semantic meaning, when comparing words using a measurement such as cosine similarity. There are also linear substructures within the word embeddings that can explain how multiple words are related to each other, making it a useful preprocessing step for natural language processing applications.

A word vector for a given word will now be defined as \mathbf{w} . Word vectors are normalised to unit length for measurement:

$$\hat{\mathbf{w}} = \frac{\mathbf{w}}{\|\mathbf{w}\|}. \quad (1)$$

All future analysis will be done using normalised word vectors, if vectors in the future are edited they will again be normalised to unit length.

2.2 Comparison of embedded words

Two words vectors \mathbf{w}_1 and \mathbf{w}_2 within a vector space can be compared by taking the dot product of their words:

$$\langle \hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2 \rangle = \sum_{i=1}^n \hat{\mathbf{w}}_{1,i} \cdot \hat{\mathbf{w}}_{2,i}. \quad (2)$$

As both word vectors are normalised, this is equivalent to the cosine similarity between the two word vectors. A cosine similarity closer to 1 means that the vectors are similar to each other, while a cosine similarity of 0 means that the vectors are orthogonal to each other.

In addition to comparisons between individual word vectors, we can compare an individual word vector to a set of word vectors. This is done by finding the mean of the set, normalizing the resulting vector and calculating the dot product with the individual word vectors as follows:

$$\langle \hat{\mathbf{w}}, \hat{\boldsymbol{\mu}} \rangle = \sum_{i=1}^n \hat{\mathbf{w}}_i \cdot \frac{\boldsymbol{\mu}_i}{\|\boldsymbol{\mu}\|}. \quad (3)$$

The resulting calculation gives us how closely an individual word is associated with a larger set of words. This association can be used to assess how closely related a given word is to different topics or concepts within the embedding space.

2.3 Removing Bias

To remove bias, first two vectors have to be identified that contain contrasting directions of the bias. These two vectors (\mathbf{w}_1 and \mathbf{w}_2) must be considered “opposite” of each other semantically, in terms of the bias that is required to be removed. The following method of debiasing is the same as presented in [2]:

$$\mathbf{w}_b = \hat{\mathbf{w}}_1 - \hat{\mathbf{w}}_2, \quad (4)$$

where the vector \mathbf{w}_b will have the direction of bias in the embedding (for example, he and she are different genders and could potentially be used to capture a gender direction).

Using this bias direction, all word vectors can now have that component removed by projecting them into a space that is orthogonal to the bias vector:

$$\mathbf{w}_\perp = \hat{\mathbf{w}} - (\hat{\mathbf{w}} \cdot \hat{\mathbf{w}}_b^T) \cdot \hat{\mathbf{w}}_b, \quad (5)$$

where \mathbf{w}_\perp is the original word vector with the biased component removed. The rank of the matrix of orthogonal projected vectors will be reduced by one in a non-trivial embedding set. These orthogonal word vectors are required to again be normalised for further analysis.

3 Experiments

In this paper, we conduct three experiments on semantic word embeddings. We first propose a new version of the Word Embedding Association Tests studied in [3] by using the LIWC lexica to systematically detect and measure the biases within the embedding, keeping the tests comparable with the same set of target words. We further extend this work using additional sets of target words, and compare sentiment across male and female names. Furthermore, we investigate gender bias in words that represent different occupations, comparing these associations with UK national employment statistics. In the last experiment, we use orthogonal projections [2] to debias our word embeddings, and measure the reduction in the biases demonstrated in the previous two experiments.

3.1 Data Description and Embedding

In all of our experiments, the first step is to obtain semantic vectors from a word embedding that we wish to analyse. We use GloVe embeddings [18], pre-trained using a window size of 10 words on a combination of Wikipedia from 2014, and the English Gigaword corpus [16], where each of the 400,000 words in the vocabulary for this embedding are represented by a 300-dimensional vector. These vectors capture, in a quantitative way, the nuanced semantics between words necessary to perform meaningful analysis of words, reflecting the semantics found in the underlying corpora used to build them.

The Wikipedia data includes the page content from all English Wikipedia pages as they appeared in 2014 when a snapshot was taken. The English Gigaword corpus is an archive of newswire text data from seven distinct international sources of English newswire covering several years up until the end of 2010 [16].

3.2 Experiment 1: LIWC Word Embedding Association Test (LIWC-WEAT)

In this experiment, we introduce the LIWC Word Embedding Association Test (LIWC-WEAT), where we measure the association between sets of target words with larger sets of words known to relate to sentiment and gender coming from the LIWC lexica [17]. We begin by using the target words from [3] which were originally used in [8], allowing us to directly compare our findings with the original WEAT.

Our approach differs from that of [3] in that while we use the same set of target words in each test, we use an expanded set of attribute words, allowing us

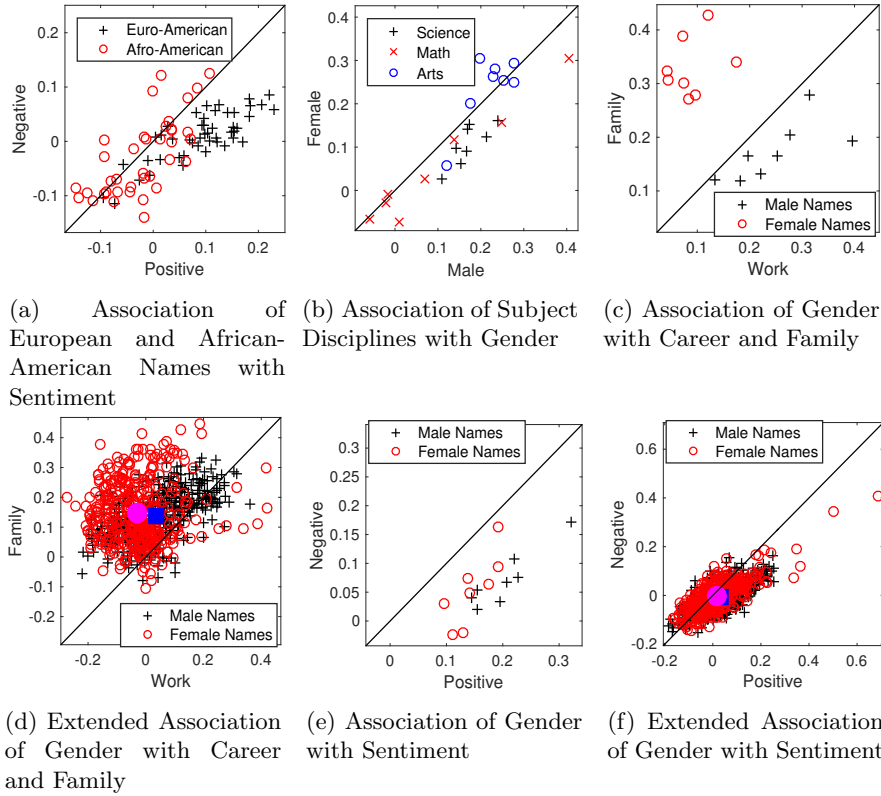


Fig. 1: Association between different words and concepts in Experiment 1, resulting from the proposed LIWC Word Embedding Association Test.

to perform a more rigorous, systematic study of the associations found within the word embeddings. For this, we use attribute words sourced from the LIWC lexica [17]. The categories specified in the LIWC lexica are based on many factors, including emotions, thinking styles, and social concerns. For each of the original word categories used in [3], we matched them with their closest equivalent within the LIWC categories, for example matching the word lists for ‘career’ and ‘family’ with the ‘work’ and ‘family’ LIWC categories.

We tested the association between each target word and the set of attribute words using the method described in Sec. 2.2, focussing on the differences in association between sentimental terms and European- and African-American names, subject disciplines to each of the genders, career and family terms with gendered names, as well as looking at the association between gender and sentiment.

Association of European and African-American Names with Sentiment Taking the list of target European-American and African-American names used in [3], we tested each of them for their associated with the positive and negative emotion concepts found in [17] by using the methodology

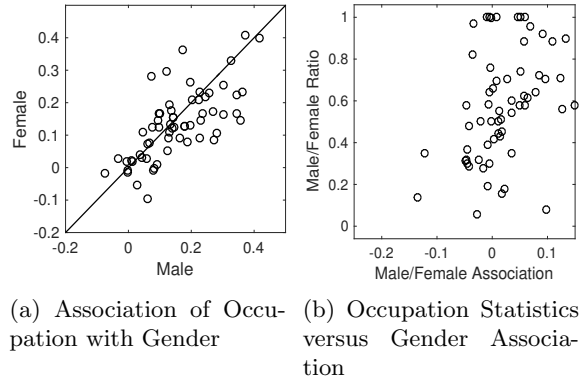


Fig. 2: Results from Experiment 2, showing the association between gender and its relation to the number of men and women working in those roles.

described by Eq. 3 in Sec. 2.2, replacing the short list of words used to originally represent pleasant and unpleasant attribute sets.

Our test found that while both European-American names and African-American names are more associated with positive emotions than negative emotions, the test showed that European-American names are more associated with positive emotions than their African-American counterparts, as shown in Fig. 1a. This finding supports the association test in [3], where they also found that European-American names were more pleasant than African-American names.

Association of Subject Disciplines with Gender A further test was conducted to find the association between words related to different subject disciplines (*e.g.* arts, maths, science) with each of the genders using the ‘he’ and ‘she’ categories from LIWC [17].

The results of our test again support the findings of [3], with Maths and Science terms being more closely associated with males, while Arts terms are more closely associated with females, as shown in Fig. 1b.

Association of Gender with Career and Family Taking the list of target gendered names used in [3], we tested each of them for their associated with the career and family concepts using the categories of ‘work’ and ‘family’ found in LIWC [17].

As shown in Fig. 1c, we found that the set of male names was more associated with the concept of work, while the female names were more associated with family, mirroring the results found in [3].

Extending this test, we generated a much larger set of male and female target names from an online list of baby names¹. Repeating the same test on this larger set of names, we found that male and female names were much less separated

¹ Baby names were taken from <http://bit.ly/2Dmqjco>, separated into two gendered lists.

Table 1: List of the top 10 occupations per gender by their association with gender.

Gender Occupations most associated with a gender

Male	Manager, Engineer, Coach, Executive, Surveyor, Secretary, Architect, Driver, Police, Caretaker, Director
Female	Housekeeper, Nurse, Therapist, Bartender, Psychologist, Designer, Pharmacist, Supervisor, Radiographer, Underwriter

than suggested by previous results, with only minor differences between the two, as shown in Fig. 1d.

Association of Gender with Sentiment Extending the number of tests performed in the original WEAT study, we additionally tested the set of target male and female names and computed their association with the positive and negative emotions. We found that both sets of names are considered to be positive, similarly to the European-American and African-American names used in the previous test, but with male names appearing to be slightly more positive, as shown in Fig. 1e.

We further tested these associations using our extended list of gendered baby names, as in Sec. 3.2, finding that there is no clear difference between the positive and negative sentiment attached to names of different gender in the word embedding.

3.3 Experiment 2: Associations between Occupations and Gender

In this experiment, we test the association between different occupations and gender categories coming from LIWC [17]. The association between each of the occupations is further contrasted against official employment statistics for the United Kingdom detailing the actual number of people working in each job role.

Association of Occupation with Gender We first generated a list of 62 occupations from data published by the Office of National Statistics [15], filtering the list to only include those occupations for which there is reliable employment statistics and can be summarised by a single word in the embedding, *e.g.* doctor, engineer, secretary. For each of these occupations, we tested their association with each of the genders, as shown in Fig. 2a, with the top ten occupations associated with each gender shown in Table 1. We found there was a 70% (p -value $< 10^{-10}$) correlation in the closeness of association between occupations and each of the gender attribute sets.

Occupation Statistics versus Occupation Association Using the list of occupations from the previous section, we compared their association with each of the genders with the ratio of the actual number of men and women working in those roles, as recorded in the official statistics [15], where 1 indicates only men

work in this role, and 0 only women. We found that there is a strong, significant correlation ($\rho = 0.57$, $p\text{-value} < 10^{-6}$) between the word embedding association between gender and occupation and the number of people of each gender in the United Kingdom working in those roles. This supports a similar finding for U.S. employment statistics using an independent set of occupations found in [3].

3.4 Experiment 3: Minimising Associations via Orthogonal Projection

In this experiment, we deploy a method for removing bias from word embeddings, first published in [2], and repeat all previous association tests related to gender reported in this paper, empirically showing the effect of bias removal on the word associations.

Finding an Orthogonal Projection for Gender To remove gender from the embedding, we first need to find a projection within the space that best encapsulates the gender differences between words. To find the best projection, we began from a list of 5 gendered pronouns in LIWC [17]. For each of the pronouns, we paired them with their gender-opposite, for example pairing “he” and “she”, “himself” and “herself” and so on. Taking the word vector from the embedding for each pronoun, we computed their difference, as described in Sec. 2.3, giving us a set of 5 potential gender projections.

Each gender projection was tested against an independent set of paired gender words sourced from WordNet [13] (containing implicit gendered words such as king and queen). After applying the gender projection to the test word-pairs, following the procedure of [2], we measured the average cosine similarity between the word-pairs. The gender projection that led to the WordNet word-pairs that are most similar (highest cosine similarity) was then selected as our gender projection, corresponding to the difference between the vectors for “himself” and “herself”.

Revised Association Tests Using the orthogonal gender projection found in the previous section, we repeated the tests from the LIWC-WEAT in Sec. 3.2 that were related to gender. This included the association of science, mathematics and the arts with gender, the association of male and females names with sentiment, work and family, and the ranking of occupations by their gender association.

In Experiment 1, we previously found that the disciplines of science and maths were more associated with male terms in the embedding, while the arts were closer to female terms. The association of each of these subject disciplines with gender after orthogonal projection was found to be more balanced, with closer to equal association for both male and female terms, shown in Fig. 3a.

Male and Females names tested in [3] showed a clear distinction in their association with work and family respectively, with our replication of the test in Sec. 3.2 finding the same results. Performing the same tests again after applying

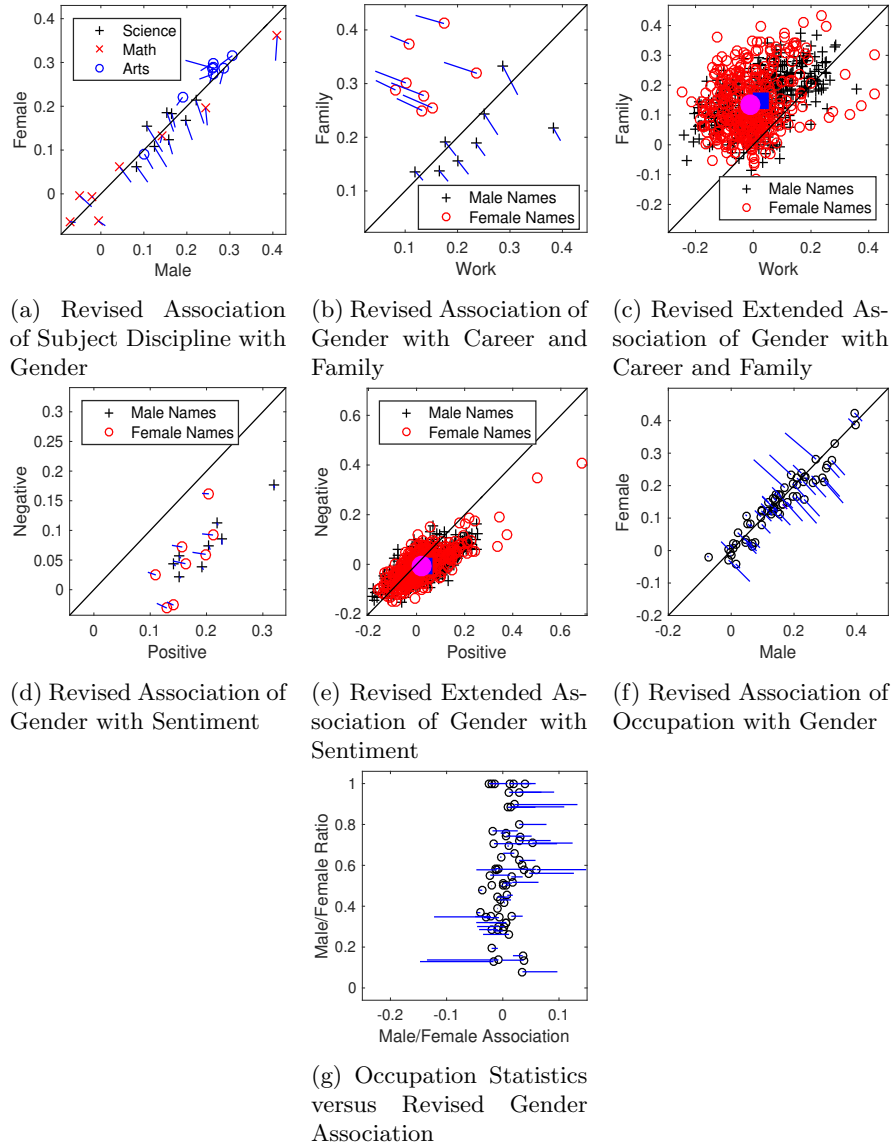


Fig. 3: Association between different words and concepts in Experiment 3 after word vectors have been debiased via orthogonal projection in the gender direction. Line-traces shown in blue indicate where points have moved from after debiasing.

the gender projection to both name lists, we wished to quantify the change in associations. We calculated the change in the distance between the centroids of each set of names before and after applying the orthogonal gender projection, finding that the association with work for males and family for females reduced,

closing the gap between male and female names by 37.5% for the target names found in the original WEAT and 66% for the extended list of names respectively.

In our experiment looking at the association of positive and negative emotions with male and female names, we found that male and female names were both positive, with male names being slightly more associated with positive emotions than female names. The same finding were also true when using a larger set of names and making the same comparison. Applying the orthogonal gender projection to the word vectors, we again looked at how much the difference between the two sets was reduced. We found that for the target names found in the original WEAT, the distance between the two sets of names was reduced by 27%, while for the extended list the difference was reduced by 40%.

In Experiment 2, we found that there was a significant correlation of 70% between the male and female association of each occupation, while comparing the associations with official statistics of the number of men and women in each role showed a correlation of 53%. Again, applying the orthogonal gender projection and repeating these tests, we found that, on average, occupations moved closer to having an equal association with each of the genders (Fig. 3f) and that their association with gender was not significantly correlated ($\rho = 0.178$, p -value = 0.167) with the number of men and women working in each role.

4 Discussion

In our experiments, we have shown the effect of one debiasing procedure for reducing the association a given word has in a word embedding generated from natural language corpora with concepts related to gender. Being able to do so relies on a set of gendered terms from which we can obtain pairings with opposite meaning, allowing us to find an orthogonal projection within the space. This will not always be possible for every type of bias that we may wish to remove (or at least reduce) in an embedding because there will not always be a suitable word vector pair that can be used to represent a given bias.

Other biases which are present may also be impossible to detect with our LIWC-WEAT method, as a pre-defined and validated list of words from LIWC were required to perform the tests. Other potentially undesired biases such as race or age are not currently able to be captured using the LIWC lexica, and thus different, carefully considered sets of words would need to be curated.

Indeed, general solutions to this problem are probably impossible, for philosophical reasons, but we believe that biases can at least be mitigated or compensated for, by removing specific subtypes of bias, given we have ways to measure and detect them in the first place. However, in this process, care should also be taken as we may introduce or compound other existing biases in the embeddings.

5 Conclusions

If we want AI to take a central position in society, we need to be able to detect and remove any source of possible discrimination, to ensure fairness and transparency, and ultimately trust in these learning systems. Principled methods to

measure biases will certainly need to play a central role in this, as will an understanding of the origins of biases, and new developments in methods that can be used to remove biases once detected.

In this paper, we have introduced the LIWC-WEAT, a set of objective tests extending the association tests in [3] by using the LIWC lexica to measure bias within word embeddings. We found bias in both the associations of gender and race, as first described in [3], while additionally finding that male names have a slightly higher positive association than female names. Biases found in the embedding were also shown to reflect biases in the real world and the media, where we found a correlation between the number of men and women in an occupation and its association with each set of male and female names. Finally, using a projection algorithm [2], we were able to reduce the gender bias shown in the embeddings, resulting in a decrease in the difference between associations for all tests based upon gender.

Further work in this direction will include removing bias in n -gram embeddings, embeddings that include multiple languages and new procedures for both generating better projections to remove a given bias, using debiased embeddings as an input to an upstream system and testing performance, and learning word embeddings which can be generated without chosen directions by construction.

Acknowledgements

AS is supported by EPSRC Centre for Communications. TLW and NC are support by the FP7 Ideas: European Research Council Grant 339365 - ThinkBIG.

References

1. Angwin, J., Larson, J., Mattu, S., Kirchner, L.: Machine bias: Theres software used across the country to predict future criminals. and its biased against blacks. ProPublica, May **23** (2016)
2. Bolukbasi, T., Chang, K.W., Zou, J.Y., Saligrama, V., Kalai, A.T.: Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: Advances in Neural Information Processing Systems. pp. 4349–4357 (2016)
3. Caliskan, A., Bryson, J.J., Narayanan, A.: Semantics derived automatically from language corpora contain human-like biases. Science **356**(6334), 183–186 (2017)
4. Cer, D., Diab, M., Agirre, E., Lopez-Gazpio, I., Specia, L.: Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. arXiv preprint arXiv:1708.00055 (2017)
5. Flaounas, I., Ali, O., Lansdall-Welfare, T., De Bie, T., Mosdell, N., Lewis, J., Cristianini, N.: Research methods in the age of digital journalism: Massive-scale automated analysis of news-contenttopics, style and gender. Digital Journalism **1**(1), 102–116 (2013)
6. Flores, A.W., Bechtel, K., Lowenkamp, C.T.: False positives, false negatives, and false analyses: A rejoinder to machine bias: There’s software used across the country to predict future criminals. and it’s biased against blacks. Fed. Probation **80**, 38 (2016)

7. Fong, R., Vedaldi, A.: Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. arXiv preprint arXiv:1801.03454 (2018)
8. Greenwald, A.G., McGhee, D.E., Schwartz, J.L.: Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology* **74**(6), 1464 (1998)
9. Jia, S., Lansdall-Welfare, T., Cristianini, N.: Freudian slips: Analysing the internal representations of a neural network from its mistakes. In: *Advances in Intelligent Data Analysis XVI*. pp. 138–148 (2017)
10. Jia, S., Lansdall-Welfare, T., Sudhahar, S., Carter, C., Cristianini, N.: Women are seen more than heard in online newspapers. *PLOS ONE* **11**(2), 1–11 (02 2016). <https://doi.org/10.1371/journal.pone.0148434>, <https://doi.org/10.1371/journal.pone.0148434>
11. Kahng, M., Andrews, P.Y., Kalro, A., Chau, D.H.P.: Activis: Visual exploration of industry-scale deep neural network models. *IEEE transactions on visualization and computer graphics* **24**(1), 88–97 (2018)
12. Lansdall-Welfare, T., Sudhahar, S., Thompson, J., Lewis, J., Team, F.N., Cristianini, N., Gregor, A., Low, B., Atkin-Wright, T., Dobson, M., et al.: Content analysis of 150 years of british periodicals. *Proceedings of the National Academy of Sciences* **114**(4), E457–E465 (2017)
13. Miller, G.A.: Wordnet: a lexical database for english. *Communications of the ACM* **38**(11), 39–41 (1995)
14. Nakov, P., Ritter, A., Rosenthal, S., Sebastiani, F., Stoyanov, V.: Semeval-2016 task 4: Sentiment analysis in twitter. In: *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*. pp. 1–18 (2016)
15. Office for National Statistics: Statistical bulletin: Annual survey of hours and earnings: 2017 provisional and 2016 revised results (2017), <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/bulletins/annualsurveyofhoursandearnings/2017provisionaland2016revisedresults>
16. Parker, R., Graff, D., Kong, J., Chen, K., Maeda, K.: English gigaword fifth edition ldc2011t07. dvd. Philadelphia: Linguistic Data Consortium (2011)
17. Pennebaker, J.W., Francis, M.E., Booth, R.J.: Linguistic inquiry and word count: Liwc 2007. Mahway: Lawrence Erlbaum Associates **71** (2001)
18. Pennington, J., Socher, R., Manning, C.: Glove: Global vectors for word representation. In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. pp. 1532–1543 (2014)
19. Ribeiro, M.T., Singh, S., Guestrin, C.: Why should i trust you?: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1135–1144 (2016)
20. Samek, W., Binder, A., Montavon, G., Lapuschkin, S., Müller, K.R.: Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems* (2017)
21. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. vol. 3, p. 6 (2017)
22. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.: Mastering the game of go with deep neural networks and tree search. *Nature* **529**(7587), 484–489 (2016)
23. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint arXiv:1703.10593 (2017)