UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Tracing the emergence of gendered language in childhood

Permalink

https://escholarship.org/uc/item/4zk7d1fx

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Prystawski, Ben
Grant, Erin
Nematzadeh, Aida
et al.

Publication Date

2020

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Tracing the emergence of gendered language in childhood

Ben Prystawski (ben.prystawski@mail.utoronto.ca)
Computer Science and Cognitive Science Programs
University of Toronto

Erin Grant (eringrant@berkeley.edu)
Department of Electrical Engineering & Computer Sciences
University of California, Berkeley

Aida Nematzadeh (nematzadeh@google.com)
DeepMind

Spike W. S. Lee (spike.lee@rotman.utoronto.ca)
Rotman School of Management and Department of Psychology
University of Toronto

Suzanne Stevenson (suzanne@cs.toronto.edu)
Department of Computer Science
University of Toronto

Yang Xu (yangxu@cs.toronto.edu)
Department of Computer Science, Cognitive Science Program
University of Toronto

Abstract

Are gender associations in general language reflected in the words spoken to and by children? Previous work has suggested that language reveals gender differences in discourse, speech style, language use and acquisition. Work in artificial intelligence has shown that word embeddings trained on large corpora reflect human gender associations. We connect this work to developmental psychology by exploring whether gender associations in word embeddings are present in the linguistic input and output of children, and if so, how early gendered language emerges. We present a computational method that quantifies the gender associations of words and use a corpus of child-caretaker speech to show that these gender associations correlate significantly with those in word embeddings. We discover that gendered word use emerges in English-speaking children around age 2, and the gender associations cannot be explained solely by variables including word length, frequency, concreteness, and valence.

Keywords: language and development; gender; child speech; word embedding; computational modelling

Introduction

Language and gender are intricately related. Work in sociolinguistics and psychology has suggested that there are gender differences in discourse, speech style, language use, and language development (Lakoff, 1973; Huttenlocher et al., 1991; Hall and Bucholtz, 2012; Holmes and Meyerhoff, 2008; Newman et al., 2008; Lovas, 2011; Coates, 2015; Laserna et al., 2014; Bamman et al., 2014). More recently, an independent line of work in computational linguistics has shown that word embeddings, or vector representations of word meaning trained on large text corpora, contain implicit gender biases such as the stereotypical association of nurses as female and engineers as male (e.g., Bolukbasi et al., 2016; Caliskan et al., 2017). Here we explore whether there is a relationship between gender associations of words said to and by children as measured by the genders of their speakers and listeners and as measured by the external metric of word embeddings (illustrated in Figure 1), and, if so, how early in life this association emerges.

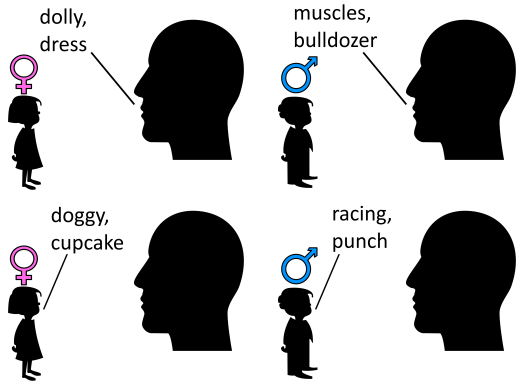


Figure 1: Illustration of gendered speech to and by children.

Gender and child language development

Research in psychology has documented gender differences in language development, but whether gender associations in general language are reflected directly in the daily speech of young children and their caretakers is an open question (Laserna et al., 2014; Bamman et al., 2014).

In fact, the literature has found no direct evidence supporting this idea. For instance, Huttenlocher et al. (1991) found that, among children between 16 and 24 months, girls' speech had a higher type-to-token ratio than boys' speech. However, the authors did not find a difference in linguistic input from caretakers that could explain this. In a meta-analysis, Leaper

©2020 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

et al. (1998) found differences in parent-to-child speech both by gender of the parent and gender of the child, such as that mothers use more supportive language when speaking to daughters than to sons. But the focus of that study is on the interaction of parent gender and child gender, not on the presence of gender differentiation in child-directed speech per se. However, there has been notable work investigating psycholinguistic correlates of gender. For example, Newman et al. (2008) found that adult men used more articles in their writing while women used more pronouns.

Evidence for gender stereotypes in language has been found in children's television programs. In particular, previous studies have analyzed the content of cartoons and children's television shows and found that language used in these shows contains gender stereotypes (Mulac et al., 1985; Aubrey and Harrison, 2004). For instance, male characters tend to associate more frequently with vocalized pauses, action verbs, and present tense verbs, whereas female characters tend to associate more with uncertainty verbs and polite terms (Mulac et al., 1985); it has also been reported that physical aggression is dominantly associated with male characters as opposed to female characters in cartoons (Luther and Legg Jr, 2010).

Although previous studies have looked at gender associations in children's media and adult-to-adult speech and writing, there is no extensive study on the presence or emergence of gender differences in the language children are exposed to or produce. We develop this line of work by exploring whether gender differentiation might be present in speech from both caretakers (i.e., linguistic input more intimate than media) and children (i.e., linguistic output) during early development.

Gender biases in word embeddings

To explore gendered language in early childhood at a comprehensive scale, we rely on an independent line of research showing that gender associations in general language are captured accurately by word embeddings. Word embeddings are vector (or distributed) representations of word meaning constructed from large-scale word co-occurrences in natural text, which have recently been used as an alternative to psycholinguistic markers of gender (Garg et al., 2018). In a typical word embedding space, words that share co-occurring context often result in close proximity, and those that share minimal context tend to be farther apart. Commonly-used methods include Word2Vec (Mikolov et al., 2013b), GloVe (Pennington et al., 2014), and fastText (Grave et al., 2018); these methods have been demonstrated to perform well on analogical reasoning tasks by representing analogies via vector algebra; for example, $\vec{\text{king}} - \vec{\text{queen}} \approx \vec{\text{man}} - \vec{\text{woman}}$.

Bolukbasi et al. (2016) found that this analogical capacity—a result of estimating semantics from word co-occurrence patterns at scale—can also reveal stereotypical gender associations in the source text corpora. For instance, they showed that the difference between the embeddings for *man* and *woman* is approximately equal to the difference between the embeddings for *computer programmer* and *homemaker*: $\vec{\text{man}} - \vec{\text{woman}} \approx \text{computer } \vec{\text{programmer}} - \vec{\text{homemaker}}$. This can be interpreted as the analogy "Man is to woman as computer programmer is to homemaker," a statement that expresses gender biases in occupation. Later work has built on these findings by showing that embedding-based gender associations are robust across source corpora and predict behavioral data in psychological tests of implicit association (Caliskan et al., 2017) and public records of gender imbalances across professions (Garg et al., 2018).

We investigate the emergence of gendered language by focusing on the linguistic input and output of children aged 1-5, from caretakers and children themselves. We explore whether language said to and by children reflects gender associations in broader culture, and if so, how early such gendered language emerges in life. We present a generic computational method that quantifies the strength of gender association in word use and test our hypothesis by examining whether gender differences measured in words from child-directed and child speech reflect gender associations in word embeddings learned from large text corpora.

Our Hypothesis

We hypothesize that the gender associations captured by word embeddings will be reflected in word usage frequency in both child-directed and child speech.

Computational methodology

Gender probability in speech

We quantify gender association of a word's usage in speech by a probability measure. This measure can be used to quantify gender differentiation either from the caretakers' perspective (i.e., child-directed speech) or children's perspective (i.e., child speech).

We measure the degree of gender association in a word based on how often that word is used toward or by children from one gender (female or male) compared to how often that word is used across both gender groups. Formally, we define the *gender probability* of a word via Bayes' rule:

$$p(g|w) \propto p(w|g)p(g) \tag{1}$$

Here g stands for gender $g \in \{f, m\}$, where f denotes female and m denotes male. w represents a word in question. We assume a uniform prior on the gender of the interlocutor (i.e., $p(f) = p(m) = 0.5$).[1] We estimate the likelihood of a word for a specific gender as the relative frequency that w is associated with that gender, e.g., the female likelihood is:

[1] To ensure that the results we report are not artifacts of our decision to use a uniform prior, we analyzed 100 random sub-samples of the corpus in which we enforced equal representation of boys and girls and ages 1 through 5 and found similar results. We used subsamples of 98,000 words for each child gender-age pair in child-directed speech and 37,000 words each in child speech. Figures 2 and 4 show the mean results across these gender-balanced subsamples.

$$p(w|f) = \frac{c(f,w)}{c(f)} \quad (2)$$

Here $c(g,w), g = f$ is the number of times word w is said to or by children with gender g, and $c(g)$ is the total count of all words said to or by children with gender g in the corpus. The (female) posterior probability is then:

$$p(f|w) = \frac{\frac{c(f,w)}{c(f)}}{\frac{c(f,w)}{c(f)} + \frac{c(m,w)}{c(m)}} \quad (3)$$

This measure accounts for the base rate difference in words said to or by f vs m. For instance, if a word w appears frequently in group f, the relative frequency of w could be quite low in group f if there are many words said to group f overall (that is, even if $c(f,w) > c(m,w)$, $p(f|w)$ could be still lower than $p(m|w)$ in principle). A word said exclusively to or by girls would have a female gender probability of 1, and a word said exclusively to or by boys would have a female gender probability of 0. We have also considered alternative measures such as log odds ratio (omitted due to space), and our results are robust to this variation.

In our analyses, we calculate gender probability separately for caretakers and children with respect to children's gender. From the caretakers' perspective, we measure the gender probability of words said to children as listeners. From the children's perspective, we measure the gender probability of words said by children (to caretakers) as speakers.

Embedding-based gender associations

Independent to our formulation of gender probability, we quantify gender associations of the public from word embeddings. We consider two representative formulations of gender association from work in computational linguistics: the Word Embedding Association Test (WEAT) (Caliskan et al., 2017) and gender Subspace Projection (Bolukbasi et al., 2016).

The Word Embedding Association Test is a common procedure for measuring associations and biases in word embeddings (Caliskan et al., 2017). In this test, the effect size of a given word's association is the difference between the mean cosine similarities between the word's vector and those of the elements of two sets of attribute words: $s(w,F,M) = \frac{\text{mean}_{f\in F}\cos(\vec{w},\vec{f}) - \text{mean}_{m\in M}\cos(\vec{w},\vec{m})}{\text{std_dev}_{x\in M\cup F}\cos(\vec{w},\vec{x})}$. Here F and M denote the two sets of attribute words—gender terms in this case, w denotes the word in question, and $\vec{x}$ denotes the vector associated with word x in the joint set of the attribute words. p-values from this test are calculated using a permutation test. We used the same sets of male and female terms used in the original word embedding association test formulated by Caliskan et al. (2017), where they found that this procedure applied to groups of target words can closely reproduce gender associations from implicit association tests in psychology (Nosek et al., 2002).

Subspace Projection (Bolukbasi et al., 2016) offers an alternative method to characterize gender association in word embeddings. This method first identifies a gender subspace by performing principal component analysis (PCA) on the vector differences between pairs of words that differ in gender, such as (woman, man) and (mother, father). The first principal component explains approximately 60% of the variance (Bolukbasi et al., 2016) and it is taken to be the axis of the one-dimensional gender subspace. An individual word's gender association is then quantified by projecting that word onto the gender subspace: $\text{proj}_G(\vec{v}) = \frac{\vec{v}\cdot\vec{b}}{|\vec{b}|}$. Here G is the gender subspace, $\vec{b}$ is the basis vector for the gender subspace identified via PCA, and $\vec{v}$ is the vector of the word in question. This returns a number between -$|\vec{v}|$ and $|\vec{v}|$, where positive numbers reflect more female-associated words and negative numbers correspond to male-associated words.

Ethayarajh et al. (2019) suggest that the WEAT overestimates the strength of associations and note that the subspace projection method is not subject to this issue. Here we consider both WEAT and subspace projection to ensure that our computational analysis is robust to methodological choices.

Data

Corpus of child and child-directed speech

We used CHILDES (MacWhinney, 2014), a large inventory of child-caretaker speech commonly used in psychology and cognitive science. We pooled data across children and caretakers from the North American section of CHILDES, filtering to include only speech from normally developing children and in naturalistic conditions. The final corpus we worked with includes 4,260,753 tokens and spans the ages of 1 to 6. Each conversation in the corpus included exactly one child, and the speakers were labelled by gender. We labelled a word as said to or by a child of a particular gender based on the gender of the child in the conversation and the participant saying the word. There were a total of 3611 children in the corpus, with 1,550 boys and 2,061 girls.[2] This is the largest publicly available corpus of child-directed speech and child speech and has been used extensively in the literature on child language development.

Pre-trained word embeddings

We used three commonly-used sets of pre-trained word embeddings: the word2vec embeddings trained on the Google News corpus (Mikolov et al., 2013a), the GloVe embeddings trained on the Common Crawl corpus, and the fastText English embeddings trained on the Common Crawl and Wikipedia corpora (Grave et al., 2018). These corpora are very large and the embeddings trained on them are commonly taken to represent general language usage and word associations present in everyday language.

[2]The percentage of words that were said to girls for each age is as follows: Age 0 - 29%, Age 1 - 49%, Age 2 - 55%, Age 3 - 44%, Age 5 - 38%, Age 6 - 9%.

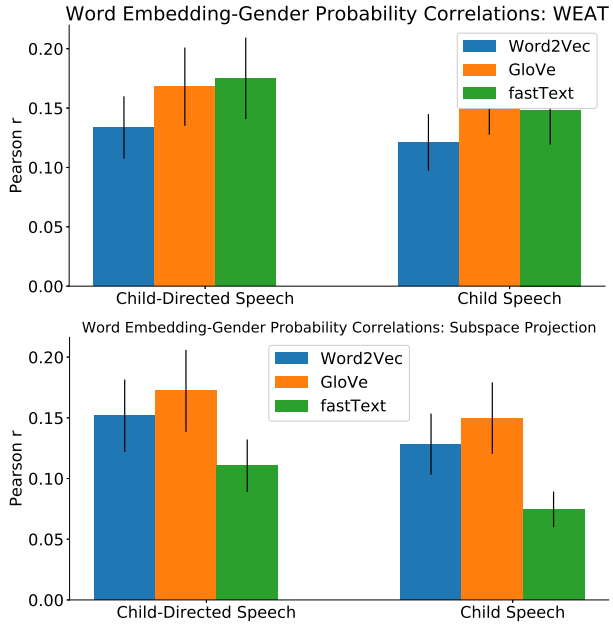


Figure 2: Correlations between gender probability and word embedding gender association. $p < 0.001$ in all cases. Analyses were based on mean association strength across subsamples of the CHILDES corpus that were balanced by age and gender. Error bars denote 95% confidence intervals.

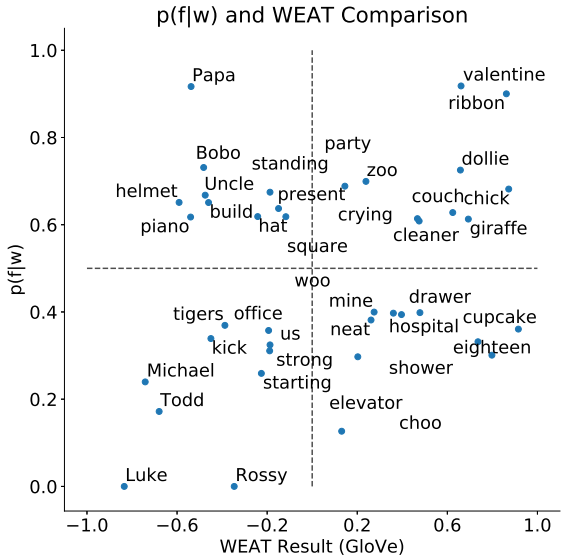


Figure 3: Word samples: Gender probability vs WEAT association.

Results

We present results in three steps. First, we show basic evidence for gendered word usage by both caretakers and children that correlate with gender associations found independently in word embeddings. Second, we show the time course of these correlations and trace the emergence of gendered language throughout child development. Third, we examine whether the gender differences found in children's linguistic environment can be explained by alternative psycholinguistic variables.

Evidence for gendered language in childhood

We start by examining whether our calculated gender probabilities of the words said to and by children in the CHILDES corpus correlate with the gender associations in everyday language estimated from the word embeddings, through the WEAT and subspace projection methods that we described. We considered all three common word embeddings.

Figure 2 summarizes the strength of correlations in these conditions.[3] We found a significant and robust Pearson correlation between gender probability and gender associations in word embeddings, with $p < 0.001$ in all of the 12 tests that we performed: 2 (child-directed vs child speech) $\times$ 2 (WEAT vs subspace) $\times$ 3 (Word2Vec, Glove, fastText). Importantly, we observed that the correlational strengths found in children's speech are similar to those found in caretakers' speech. These results provide evidence that gender associations are reflected directly in word usage during early child development both in input and output—a finding that goes beyond existing research showing the presence of gender differences in language from external media, but not from interlocutors (Mulac et al., 1985; Luther and Legg Jr, 2010).

Figure 3 shows a subset of words (in the CHILDES lexicon) illustrating the similarities and differences between gender probability in children's linguistic environment and gender associations in word embeddings from the WEAT. Words that fall in the bottom-left and top-right quadrants correspond to concordant cases between the two measures. In particular, words of action and strength (cf. Mulac et al. (1985)) such as *kick*, *strong*, and words of typical male names such as *Michael*, *Todd*, are found to be consistently male-associated based on our measure and WEAT; words related to animals such as *zoo*, *chick*, *giraffe*, and words related to female characters/toys such as *dollie*, are found to be consistently female-associated. Not all words are correlated between the two measures. For instance, male kin terms such as *Papa*, *Uncle* are more female-oriented in children's linguistic environment but more male-oriented in the word embeddings, and words such as *cupcake* and *drawer* are more gender-neutral in child development compared to the word embeddings. We intentionally sampled an equal number of words from each quadrant to illustrate both concordant and discordant cases. In a random sample of words, we would expect to see more in the top-right and bottom-left.

[3] For these analyses, we kept all words which occurred at least 50 times in the CHILDES corpus that were also represented in the vocabulary of the pre-trained word embeddings. This resulted in 2,041 words with word2vec, 2,051 words with GloVe, and 2,055 words with fastText. In the downsampled version, we used a threshold of 20 occurrences to account for the smaller corpus size. There were on average 1316.46 words for word2vec, 1328.11 for GloVe, and 1324.89 for fastText.

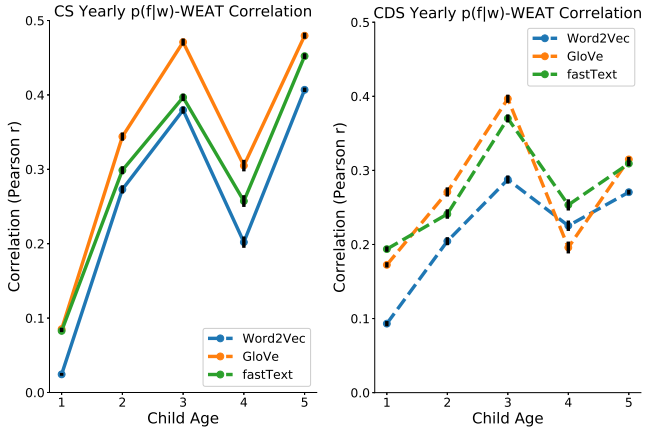


Figure 4: Developmental time course of correlation between gender probability in child (left) and child-directed (right) speech and gender association in word embeddings. The correlation was statistically insignificant in child speech in year 1 with GloVe and Word2Vec and significant for all other ages and years. Error bars denote 95% confidence intervals run on 100 sub-samples of the corpus, balanced for gender and age.

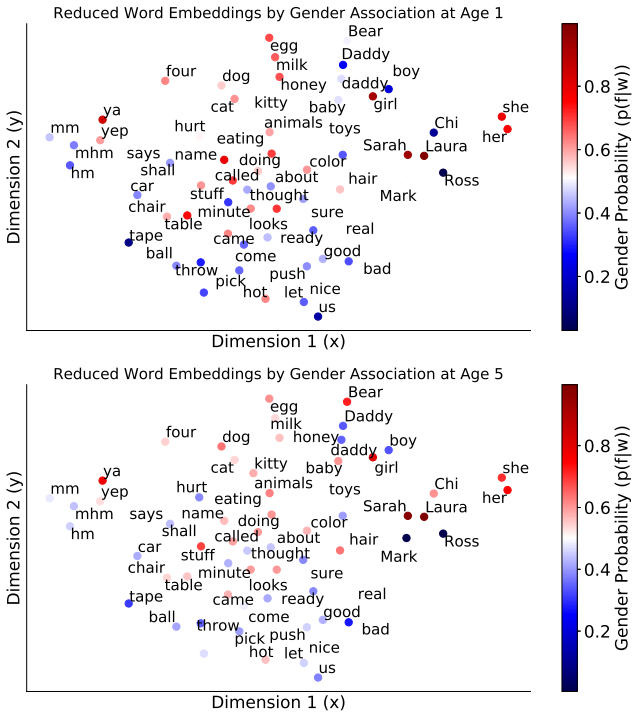


Figure 5: Visualization of words in child-directed speech that show high gender probabilities, for age groups 1 and 5 in development. Semantic space is constructed from dimensionality-reduced fastText word embeddings. Dimension 1 captures some aspect of animacy, while Dimension 2 captures some of the variance along the gender axis. Colorbar indicates the scale of gender probability, with 1 indicating words exclusively uttered by girls and 0 exclusively by boys.

Developmental time course of gendered language

To examine the time course of gendered language, we separated the CHILDES data by year. We focused on tracking the emergence of gendered language for children ages 1 through 5 since data for ages 0 and 6 is relatively sparse.[4] For each year, we measured correlations between word embedding associations and words said at least 25 times by and to children of that age in the CHILDES corpus. We varied the frequency threshold between 10 and 50 and the results are robust.

We computed gender probability for the subset of words in the corpus said by children of each age, then measured the correlation between the gender probabilities and WEAT associations. Figure 4 shows that the correlation strength between word embedding associations and gender probability in child speech gradually increases from age 1 to 5, with a dip at age 4.[5] The correlation at age 1 is not significant ($p = 0.081$ for GloVe) but the correlations for all other ages are highly significant with $p < 0.001$. The trend in child-directed speech is comparatively flatter. The correlation is significant from age 1 and increases more gradually than in child speech. These results suggest not only that gendered word uses in children's linguistic environment correlate with gender associations in word embeddings, but that they also emerge very early in life—around the age of 2 in child speech.

Figure 5 visualizes gender probability of a sample of words in child speech in a 2-D word-embedding space via t-SNE dimensionality reduction (Maaten and Hinton, 2008). We focused on the 30 words with the highest aggregate gender probability and the 30 with the lowest among words that occur at least 500 times in the corpus. Gender probability follows a color scale, where red corresponds to words said more to girls and blue corresponds to words said more to boys. We found clusters of words in the embedding space that share similar gender probabilities belonging to each gender group. For instance, morally valenced words (cf. Luther and Legg Jr (2010)) such as *good*, *bad*, and *nice* are close together in embedding space and all said more to boys, while animal words such as *animals*, *cat*, and *dog* are close to each other and said more to girls. We also found this distinction between male and female-oriented words to be persistent through the developmental course, illustrated in children at ages 1 and 5.

Psycholinguistic correlates of gender probability

In our final analyses, we investigated the degree to which our measure of gender probability correlates with alternative psycholinguistic variables that are representative in the literature of language development. If our findings so far reflect gender-specific aspects of language development, then we expect that gender probabilities should correlate with word embedding associations beyond other confounding variables that

[4]The number of tokens per age in the CHILDES corpus is as follows: Age 0 - 356,789, Age 1 - 929,601, Age 2 - 945,358, Age 3 - 654,719, Age 4 - 945,501, Age 5- 393,505, Age 6 - 167,965.

[5]This figure shows results from using WEAT to measure associations in word embeddings, but similar results occur when using the subspace projection method.

Table 1: Correlations of gender probability and psycholinguistic metrics in child-directed (CDS) and child (CS) speech.

Variable correlation	Pearson ρ	p
Length (CDS)	-0.02	0.006
Length (CS)	0.0063	0.41
Log-frequency (CDS)	0.10	< 0.001
Log-frequency (CS)	0.13	< 0.001
Concreteness (CDS)	0.060	< 0.001
Concreteness (CS)	0.080	< 0.001
Valence (CDS)	0.062	< 0.001
Valence (CS)	0.10	< 0.001

could explain the gender effect. We considered four word metrics: form length, usage frequency, concreteness, and valence. We computed the length of each word and estimated the frequency in CDS and CS directly from the CHILDES data by counting the occurrences of the word said to or by all children respectively. We took concreteness and valence ratings of words from existing large-scale behavioral experiments in Warriner et al. (2013) and Brysbaert et al. (2014). In these experiments, humans rated words' concreteness or valence from 1-10. Ratings were averaged across participants. We then measured the correlation between gender probabilities in CDS and CS with fastText WEAT associations and each of the four variables.

Table 1 summarizes the Pearson correlation coefficients and p-values from our analyses. We found significant but small ($\rho \leq 0.13$) correlations between the gender probability of a word and the metrics that we considered. We found that words with shorter length, higher frequency, more positive valence, and higher concreteness tend also to be said more to and by girls than boys. The correlation between gender probability in child-directed speech and valence is consistent with the previous finding from Leaper et al. (1998) that mothers use more supportive language when speaking to girls compared with boys. In addition, the correlation between gender probability in child speech and word concreteness is also consistent with the result from Newman et al. (2008) that adult women use more concrete vocabulary than adult men on average. We next applied linear regression to all four psycholinguistic variables to account for gender probability in both child-directed speech and child speech. The coefficients with 95% confidence intervals are summarized in Table 2. These variables explain the variance in gender probability quite well, as the R^2 values were 0.602 for CDS and 0.598 for CS. Finally, we analyzed the partial correlation between gender probability and gender WEAT results while controlling for the four psycholinguistic variables. We focused on the 955 words for which this data was available in the datasets from Warriner et al. (2013) and Brysbaert et al. (2014). Results of this analysis are summarized in Table 3. Controlling for these factors only reduces the correlation strength by 0.04-0.05 ($p < 0.001$ in all cases). Taken together, our results

Table 2: Coefficients from linear regression using psycholinguistic correlates of gender probability in CDS and CS.

Variable	CS	CDS
Length	0.0229 [-0.050, 0.096]	0.244 [0.185, 0.302]
Log-freq.	0.238 [0.185, 0.290]	0.249 [0.201, 0.297]
Concrete.	0.23 [0.188, 0.275]	0.169 [0.135, 0.204]
Valence	0.245 [0.185, 0.304]	0.222 [0.173, 0.270]

Table 3: Full and partial correlations between word embedding associations and gender probability.

Embedding	CS	CS (partial)	CDS	CDS (partial)
Word2Vec	0.282	0.236	0.254	0.193
GloVe	0.377	0.341	0.294	0.24
fastText	0.307	0.264	0.295	0.241

show that the gender probabilities in child development are complementary to other factors, as the variability in gender effects cannot be explained solely by confounding variables.

Discussion

We have presented to our knowledge the first large-scale computational investigation that shows evidence for gender differences in the linguistic input and output of children. Our emphasis is on understanding how gendered language emerges early in life and whether it mirrors the public's gender associations, which we approximated through established word embedding methods. Our results indicate a significant correlation between gender probability in child/child-directed speech and gender associations in word embeddings. We have shown that this correlation appears as early as age 2 in child speech and persists throughout development. We have demonstrated that these gender effects in language development are complementary to the psycholinguistic variables of word length, frequency, concreteness, and valence.

Our findings are based on the analyses of English-speaking children and caretakers. Although corpora in other languages are available in CHILDES, the scales of those corpora are smaller than in English. Future work should explore the generality of our findings to other languages and compare how gender differentiation might be reflected in speech beyond childhood.

Extending the rich literature on language development and gender, our findings suggest that social biases may exert an early influence on child development, both in the input (caretakers' speech) and the output (children's speech) of language in childhood.

Acknowledgments

This work was supported by a Research Grant from the Institute for Gender and the Economy. We thank Claire Bowern for constructive comments and Michael Frank for helpful pointers. BP is funded by an NSERC USRA award. YX is funded through an NSERC Discovery Grant, a SSHRC Insight Grant, and a Connaught New Researcher Award. SWSL is funded by an Early Researcher Award (Ontario Ministry of Research, Innovation and Science) and an Insight Grant (SSHRC). SS is funded through NSERC grant RGPIN-2017-06506.

References

Aubrey, J. S. and Harrison, K. (2004). The gender-role content of children's favorite television programs and its links to their gender-related perceptions. *Media Psychology*, 6:111–146.

Bamman, D., Eisenstein, J., and Schnoebelen, T. (2014). Gender identity and lexical variation in social media. *Journal of Sociolinguistics*, 18(2):135–160.

Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *NeurIPS*, pages 4349–4357.

Brysbaert, M., Warriner, A. B., and Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known english word lemmas. *Behavior Research Methods*, 46(3):904–911.

Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Coates, J. (2015). *Women, men and language: A sociolinguistic account of gender differences in language*. Routledge.

Ethayarajh, K., Duvenaud, D., and Hirst, G. (2019). Understanding undesirable word embedding associations. *ArXiv Preprint arXiv:1908.06361*.

Garg, N., Schiebinger, L., Jurafsky, D., and Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.

Grave, E., Bojanowski, P., Gupta, P., Joulin, A., and Mikolov, T. (2018). Learning word vectors for 157 languages. *arXiv preprint arXiv:1802.06893*.

Hall, K. and Bucholtz, M. (2012). *Gender articulated: Language and the socially constructed self*. Routledge.

Holmes, J. and Meyerhoff, M. (2008). *The handbook of language and gender*. John Wiley & Sons.

Huttenlocher, J., Haight, W., Bryk, A., Seltzer, M., and Lyons, T. (1991). Early vocabulary growth: relation to language input and gender. *Developmental Psychology*, 27(2):236.

Lakoff, R. (1973). Language and woman's place. *Language in society*, 2(1):45–79.

Laserna, C. M., Seih, Y.-T., and Pennebaker, J. W. (2014). Um... who like says you know: Filler word use as a function of age, gender, and personality. *Journal of Language and Social Psychology*, 33(3):328–338.

Leaper, C., Anderson, K. J., and Sanders, P. (1998). Moderators of gender effects on parents' talk to their children: A meta-analysis. *Developmental Psychology*, 34(1):3.

Lovas, G. S. (2011). Gender and patterns of language development in mother-toddler and father-toddler dyads. *First language*, 31(1):83–108.

Luther, C. A. and Legg Jr, J. R. (2010). Gender differences in depictions of social and physical aggression in children's television cartoons in the us. *Journal of Children and Media*, 4:191–205.

Maaten, L. v. d. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605.

MacWhinney, B. (2014). *The CHILDES project: Tools for analyzing talk, Volume II: The database*. Psychology Press.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In *NeurIPS*, pages 3111–3119.

Mulac, A., Bradac, J. J., and Mann, S. K. (1985). Male/female language differences and attributional consequences in children's television. *Human Communication Research*, 11:481–506.

Newman, M. L., Groom, C. J., Handelman, L. D., and Pennebaker, J. W. (2008). Gender differences in language use: An analysis of 14,000 text samples. *Discourse Processes*, 45(3):211–236.

Nosek, B. A., Banaji, M. R., and Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice*, 6(1):101.

Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.

Warriner, A. B., Kuperman, V., and Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior Research Methods*, 45(4):1191–1207.