



**HAL**  
open science

# Debiasing Gender Bias in Information Retrieval Models

Dhanasekar Sundararaman, Vivek Subramanian

► **To cite this version:**

Dhanasekar Sundararaman, Vivek Subramanian. Debiasing Gender Bias in Information Retrieval Models. 2022. ⟨hal-03797552⟩

**HAL Id: hal-03797552**

**<https://hal.science/hal-03797552v1>**

Preprint submitted on 4 Oct 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

# Debiasing Gender Bias in Information Retrieval Models

Dhanasekar Sundararaman  
dhanasekar.sundararaman@duke.edu  
Duke University  
Durham, North Carolina, USA

Vivek Subramanian  
viveksub@amazon.com  
Amazon Alexa AI  
Boston, Massachusetts, USA

## ABSTRACT

Biases in culture, gender, ethnicity, etc. have existed for decades and have affected many areas of human social interaction. These biases have been shown to impact machine learning (ML) models, and for natural language processing (NLP), this can have severe consequences for downstream tasks. Mitigating gender bias in information retrieval (IR) is important to avoid propagating stereotypes. In this work, we employ a dataset consisting of two components: (1) relevance of a document to a query and (2) “gender” of a document, in which pronouns are replaced by male, female, and neutral conjugations. We definitively show that pre-trained models for IR do not perform well in zero-shot retrieval tasks when full fine-tuning of a large pre-trained BERT encoder is performed and that lightweight fine-tuning performed with adapter networks improves zero-shot retrieval performance almost by 20% over baseline. We also illustrate that pre-trained models have gender biases that result in retrieved articles tending to be more often male than female. We overcome this by introducing a debiasing technique that penalizes the model when it prefers male over female, resulting in an effective model that retrieves articles in a balanced fashion across genders.

## KEYWORDS

Gender bias, bias in retrieval models, zero-shot retrieval

## 1 INTRODUCTION

Differences in culture, gender, ethnicity, etc. naturally lead to conscious and unconscious bias among humans [15]. These biases have existed for decades and have affected many areas of human social interaction such as admissions to higher institutions, job selection, and financial compensation. Such human biases are incorporated into artificial intelligence (AI) models, which are trained on examples generated by humans. It is critical to identify and mitigate these biases as AIs become more ubiquitous in day to day tasks.

Biases have been shown to impact all types of machine learning (ML) models, including computer vision (CV) and natural language processing (NLP) [19, 20]. For instance, algorithms for facial recognition have been shown to disproportionately identify black and hispanic males as criminals among a diverse set of images [13]. In addition, it was shown that algorithms developed by leaders in facial recognition more often misclassified black women than white males [10]. In the NLP community, embeddings trained on large corpora of data [14] were more likely to associate African American names with unpleasant words [4]. Generative models for images from text, when queried to generate images of doctors or engineers, typically generate male rather than female individuals [5]. These biases can feed into various downstream NLP tasks [18, 21] including summarization, generation [17], and information retrieval.

Information retrieval (IR) is the task of identifying the document in a database which is most relevant to a given query. There are two main tasks in IR: (1) representing queries and documents and (2) scoring relevance between the two. With the advent of deep learning, task (1) is typically accomplished by utilizing large, pre-trained language models such as BERT to generate embeddings for sequences of tokens [6]. Then, scoring relevance can be accomplished by computing some similarity metric (e.g., cosine similarity [16]) between query and document embeddings. One of the main shortcomings of BERT is that it is trained on two large corpora of text – BooksCorpus (800M words) and Wikipedia (2500M words) [6] – which are known to suffer from gender bias [1]. The bias present in BERT could be realized through its performance on a number of downstream tasks ranging from classification to tagging.

Eliminating gender bias in IR is important since biased retrieval can result in unwanted propagation of negative gender stereotypes. For instance, as shown in [7], word embeddings derived from over 100 years of text data evolve with occupational patterns observed among males and females. Utilizing these embeddings, which associate gender with occupational roles that may not be relevant in today’s context, can bias an IR model. Models pretrained using such embeddings tend to be biased towards a (gender, occupation) pair even assuming there is an ideal dataset in which the frequency of each gender label is equal. For example, an IR model could reinforce stereotypes of traditional female roles such as secretary, nurse, or housekeeper while diminishing the roles of women who are engineers, mechanics, or carpenters [7].

Very few methods have been proposed to reduce gender bias in IR. In image retrieval, it has been shown that women are less likely to be identified as working at a computer [8]. This bias was mitigated by adding two complementary loss functions which both confuse the model when gender information is not present in an image and make the model more confident when it is. Similarly, in text retrieval, methods have focused on balancing the number of examples associated with each gender, but this does not necessarily eliminate more complex biases at a model-level that correlate gender with certain features [2].

In addition, with the plethora of NLP/IR datasets that are being collected and released, it becomes increasingly unlikely that the distribution of the data on which a model must be evaluated matches that of the data used for pretraining. This motivates the need for zero-shot learning and retrieval methods, which allow a model to generalize to unseen dataset distributions. In addition, evaluating a pretrained model on a zero-shot basis can help us to learn more about the bias of a model since the training and test distributions are no longer connected.

In our experiments, we employ the GrepBiasIR dataset [9], which consists of seven categories of query-document pairs. (More details provided in section 3.1.) For a very small dataset, BERT,

**Table 1: Zero-shot retrieval results using full fine-tuning.**

Train/Test	Sex & Relationship	Career	Child Care	Appearance	Cognitive Capabilities	Domestic Work	Physical Capabilities	Average*
Sex & Relationships	0.9565	0.4833	0.4167	0.5238	0.5139	0.5000	0.4912	0.4882
Career	0.5290	0.8250	0.6548	0.5476	0.6111	0.4889	0.4912	0.5538
Child Care	0.5942	0.5250	1.0000	0.5119	0.5000	0.4111	0.5526	0.5158
Appearance	0.4783	0.4250	0.5476	1.0000	0.4028	0.5444	0.5263	0.4874
Cognitive Capabilities	0.5217	0.4917	0.5000	0.4762	0.8333	0.5111	0.4737	0.4957
Domestic Work	0.5435	0.5000	0.4643	0.6071	0.5972	0.9000	0.4737	0.531
Physical Capabilities	0.4855	0.4750	0.5476	0.5476	0.5833	0.5333	1.0000	0.5287

being a model with millions of parameters, is likely to overfit with full fine-tuning. Thus, there is a need for lightweight fine-tuning that effectively uses BERT parameters while also generalizing well. Adapter networks, single layer networks that can be introduced into a complex model, achieve this by allowing us to freeze BERT parameters while updating adapter layers. We show that this form of fine-tuning performs better on GrepBiasIR than the full fine-tuning approach, producing an average improvement of over 20%.

To summarize our contributions, in this paper, we analyze the information retrieval capability of a BERT model over several categories of a bias IR dataset. By fine-tuning on one category and testing on the others, we are able to show the bias of BERT towards different genders at a more granular level. It is imperative that fine-tuning takes place only on a single category of data because we want to show that our method generalizes well to classes outside the one on which the model is fine-tuned. We also introduce a regularization-based debiasing technique for mitigating the aforementioned gender bias. Our methods can be used to debias BERT for a wide variety of applications, including but not limited to IR.

## 2 RELATED WORKS

The topic of gender bias in ML models is highly debated and studied. While there have been some significant contributions in the past, not a lot of emphasis has been shown particularly on retrieval models, let alone zero-shot transfer. Works such as [3] identified the gender bias problem in ML models, showing the intrinsic bias in word embeddings that could potentially propagate to numerous downstream tasks. [22] proposed methods for gender-neutral static embeddings by preserving gender information in certain dimensions of a word embedding while others are free from gender information.

In retrieval, [12] has methodologically analyzed gender bias in prominent retrieval models. They found that all models exhibited strong bias towards male and that the bias is pronounced in neural models, which corroborates our findings. The neural models exhibit significant bias because of the pre-trained embeddings trained from biased corpora. To the best of our knowledge, we are the first to identify gender bias in retrieval models on a zero-shot basis, and to propose methods to mitigate them while not compromising retrieval performance.

## 3 PRELIMINARIES

### 3.1 Datasets

For our experiments, we employ the GrepBiasIR dataset [9], which consists of 819 query, document title, and document content tuples divided roughly evenly into seven categories: Appearance, Career, Child Care, Cognitive Capability, Domestic Work, Physical Capability, and Sex and Relationships. Queries in each category are each paired with two documents: one which is relevant to the query and one which is not. In addition, three versions of each document have been created by substituting pronouns with their male ( $M$ ), female ( $F$ ), and neutral ( $N$ ) forms. This allows us to assess retrieval performance, bias of a pretrained model, and bias correction after fine-tuning.

### 3.2 Training

For full fine-tuning, we train for 2 epochs using an AdamW Optimizer with learning rate  $2 \times 10^{-5}$  and batch size  $B$  of 8. This is the maximum number of epochs on which the model can be trained without validation accuracy saturating. The model is trained with cross-entropy loss. For adapter fine-tuning, we train for 8 epochs with learning rate  $1 \times 10^{-4}$  using the same optimizer and batch size.

## 4 METHODS

### 4.1 Adapter-based fine-tuning

Our dataset consists of  $N$  query-document pairs  $x_i \forall i \in \{1, \dots, N\}$ . To classify a document’s relevancy to the query, we first employ a BERT encoder which maps each data point to an embedding  $x_i \in \mathbb{R}^D$ , where  $D = 768$ , via the BERT CLS token [6]. This query-document embedding is then directly passed through a binary classifier, which is the adapter layer in our model, to obtain a score for the data point. Mathematically, we have:

$$y_i = \sigma(\mathbf{A}x_i) = \sigma(z_i) \tag{1}$$

where  $\mathbf{A} \in \mathbb{R}^D$  is a learned vector of weights,  $z_i \in \mathbb{R}$  are logits for relevancy scores, and  $\sigma(\cdot) : \mathbb{R} \mapsto (0, 1)$  is the logistic sigmoid activation function:  $\sigma(z) = \frac{1}{1 + \exp(-z)}$ . We fine-tune these weights on a per-category basis to obtain retrieval scores for all other categories. For each data point, the loss is the binary cross-entropy loss:

$$\mathcal{L}_i = -t_i \log(y_i) - (1 - t_i) \log(1 - y_i) \tag{2}$$

where  $t_i \in \{0, 1\}$  is the target label for data point  $x_i$ . This is averaged over a batch, and weights in the output adapter layer are updated by gradient descent. Importantly, the weights of the BERT encoder

**Table 2: Zero-shot retrieval using adapter-based fine-tuning.**

Train/Test	Sex & Relationships	Career	Child Care	Appearance	Cognitive Capabilities	Domestic Work	Physical Capabilities	Average*
Sex & Relationships	0.9565	0.6250	0.5357	0.6548	0.8194	0.7222	0.7193	0.6794
Career	0.8188	0.8917	0.5714	0.7024	0.7500	0.7333	0.7018	0.7130
Child Care	0.7536	0.6833	0.9048	0.7500	0.6944	0.7222	0.7105	0.7190
Appearance	0.7826	0.6167	0.5714	0.9286	0.7917	0.7222	0.7105	0.6992
Cognitive Capabilities	0.8188	0.6083	0.5357	0.6429	0.9306	0.6778	0.6842	0.6613
Domestic Work	0.7899	0.6167	0.5476	0.6667	0.7917	0.9000	0.7018	0.6857
Physical Capabilities	0.8188	0.5750	0.6071	0.7143	0.7778	0.6889	0.8947	0.6970

**Table 3: Fraction of male, female, and neutral articles with highest logits. Model is biased towards male.**

Train/Test	Sex & Relationships			Career			Child Care			Appearance			Cognitive Capabilities			Domestic Work			Physical Capabilities			Average		
	M	F	N	M	F	N	M	F	N	M	F	N	M	F	N	M	F	N	M	F	N	M	F	N
Sex & Relationships	0.13	0.18	0.69	0.43	0	0.57	0.45	0	0.55	0.25	0.33	0.42	0.14	0	0.86	0.33	0.33	0.33	0.46	0.23	0.31	0.32	0.2	0.58
Career	0.14	0.14	0.72	0.50	0.06	0.44	0.40	0.20	0.40	0.33	0.17	0.50	0	0	1	0.40	0.20	0.40	0.36	0.18	0.46	0.31	0.13	0.56
Child Care	0.08	0.25	0.67	0.15	0.31	0.54	0.25	0.17	0.58	0.08	0.50	0.42	0	0.14	0.86	0.40	0.20	0.40	0.21	0.50	0.29	0.16	0.3	0.53
Appearance	0.14	0.22	0.64	0.23	0.23	0.54	0.50	0	0.50	0.25	0.25	0.50	0	0	1	0	0.20	0.80	0.36	0.45	0.19	0.21	0.19	0.6
Cognitive Capabilities	0.13	0.20	0.67	0.33	0.20	0.47	0.42	0	0.58	0.25	0.17	0.58	0.25	0	0.75	0.17	0	0.83	0.15	0.23	0.62	0.25	0.11	0.64
Domestic Work	0.13	0.20	0.67	0.50	0.21	0.29	0.50	0	0.50	0.25	0.25	0.50	0	0	1	0.14	0.43	0.43	0.46	0.46	0.08	0.28	0.22	0.5
Physical Capabilities	0.20	0.13	0.67	0.31	0.31	0.38	0.40	0	0.60	0.25	0.25	0.50	0.29	0	0.71	0.40	0.20	0.40	0.31	0.31	0.38	0.30	0.18	0.52

remain fixed throughout the fine-tuning process, so each update is roughly  $\mathcal{O}(D)$ .

## 4.2 Mitigating gender bias

While the above method is an effective way to build a classifier for determining relevance of a document to a query, in practice, we find that documents whose pronouns are male-gendered tend to be more relevant to the same query than the corresponding female-gendered documents. This reflects the well-known gender bias in the BERT model.

To mitigate this bias, we employ a regularization term which penalizes for preference of one gender over another. In each batch (indexed by  $j$ ), we randomly sample without replacement pairs (indexed by  $k$ ) of data points  $(x_k, x'_k)$  whose documents are of different genders  $(g_k, g'_k)$ , where  $g_k, g'_k \in \{M, F, N\}$ . We then compute the sums of squared differences between relevancy logits  $(z_k, z'_k)$  of query-document pairs within a batch that are of different genders. These are averaged across all samples in a batch to compute the regularization term  $\mathcal{R}_j$  for that batch. Each pair of differences is weighted by an  $\alpha_k \in \mathbb{R}^+$ , whose value depends on the genders of the documents involved in the pair. Thus, there are a total of  $\binom{G}{2} = 3$  unique  $\alpha_k$ s, where  $G = 3$  is the number of genders, which are tuned as hyperparameters by grid search over 0 to 2 in increments of 0.25. Mathematically, we have:

$$\mathcal{R}_j = \frac{1}{N_j} \sum_{\{k : g_k \neq g'_k\}} \alpha_k (\mathcal{G}(x_k), \mathcal{G}(x'_k))(z_k - z'_k)^2 \quad (3)$$

where  $N_j$  is the total number of pairs in batch  $j$  and  $\mathcal{G}(x_k) : x_k \mapsto g_k$ . Thus, the total loss for a batch becomes:

$$\mathcal{L}_j = \frac{1}{B} \sum_{i=1}^B \mathcal{L}_i + \mathcal{R}_j \quad (4)$$

**Table 4: Top words in each category based on TF-IDF.**

Category	Top-3 TF-IDF words
Sex and Relationships	love (0.18), emotional (0.06), friend (0.06)
Career	job (0.10), career (0.08), home (0.08)
Child Care	baby (0.34), parents (0.11), children (0.08)
Appearance	beauty (0.10), hair (0.10), attractive (0.06)
Cognitive Capabilities	iq (0.09), intelligence (0.06), gifted (0.06)
Domestic Work	home (0.14), cleaning (0.11), house (0.10)
Physical Capabilities	athletes (0.17), fat (0.11), weight (0.10)

## 5 RESULTS AND DISCUSSION

Table 1 shows the retrieval accuracy obtained from our full fine-tuning experiments on each of the seven categories. (In all tables, columns denoted Average\* indicate that the average was taken across a row without the diagonal element, to quantify zero-shot performance.) Baseline accuracy on each class is roughly chance (50%) before fine-tuning. After full fine-tuning of BERT, in which all layers of BERT are fine-tuned, we find that the retrieval performance on data from the same class as that in which fine-tuning was performed greatly improves (retrieval accuracy of  $0.9307 \pm 0.0782$ ). For the remaining classes, we find that zero-shot retrieval performance still hovers around chance (average of averages for each class  $0.5144 \pm 0.0252$ ), illustrating the poor generalization performance of full fine-tuning.

Table 2 shows that fine-tuning using an adapter layer significantly boosts performance compared to full fine-tuning. Specifically, average performance on the same class as the fine-tuning data is  $0.9153 \pm 0.0239$ . Although the average performance is slightly lower, the error bar is much tighter, likely because there are fewer parameters over which additional uncertainty is introduced. Average of

**Table 5: Zero-shot retrieval after de-biasing.**

Train/Test	Sex & Relationships	Career	Child Care	Appearance	Cognitive Capabilities	Domestic Work	Physical Capabilities	Average*
Sex & Relationships	0.7681	0.5750	0.5238	0.6310	0.7639	0.6778	0.6579	0.6382
Career	0.6594	0.6416	0.6785	0.6904	0.7083	0.6888	0.6842	0.6849
Child Care	0.7391	0.6167	0.6667	0.7500	0.7500	0.7222	0.6842	0.7103
Appearance	0.7681	0.6250	0.6429	0.7857	0.7917	0.7222	0.7105	0.7100
Cognitive Capabilities	0.6231	0.4833	0.5595	0.5833	0.8472	0.5888	0.6754	0.5856
Domestic Work	0.7899	0.6167	0.5476	0.6667	0.7917	0.9000	0.7018	0.6857
Physical Capabilities	0.7463	0.6083	0.6309	0.7261	0.8333	0.7555	0.7192	0.7167

**Table 6: Percentage difference between average male and female, before and after de-biasing.**

	Average Before			Average After			M - F	
	M	F	N	M	F	N	Before	After
Sex & Relationships	0.32	0.2	0.58	0.21	0.27	0.52	0.12	<b>0.06</b>
Career	0.31	0.13	0.56	0.28	0.22	0.5	0.18	<b>0.06</b>
Child Care	0.16	0.3	0.53	0.22	0.26	0.52	0.14	<b>0.04</b>
Appearance	0.21	0.19	0.6	0.21	0.24	0.55	0.02	0.03
Cognitive Capabilities	0.25	0.11	0.64	0.27	0.36	0.37	0.14	<b>0.09</b>
Domestic Work	0.28	0.22	0.5	0.28	0.22	0.5	0.06	0.06
Physical Capabilities	0.3	0.18	0.52	0.31	0.25	0.44	0.12	<b>0.06</b>

averages for fine-tuning with an adapter model on each class is  $0.6935 \pm 0.0199$ , an improvement of almost 20%, indicating much greater zero-shot generalizability after this type of fine-tuning.

Table 3 shows a detailed breakdown by gender of the results from Table 2. Specifically, for each class on which we fine-tuned, we show, for the correctly retrieved examples, the fraction for which the gender label with the highest relevancy logit value was  $M$ ,  $F$ , or  $N$ . On average, we find that the absolute difference between  $M$  and  $F$  fractions is  $11.14 \pm 5.39\%$ , exemplifying the strong gender bias of the BERT model. Note that examples for which the model retrieved the male and female versions of a document with equal probability were discarded to avoid diluting the bias with examples for which the model was equally confident about both genders.

We find that, on average, in six out of the seven categories, the male version of a document is preferred over a female version, with the female only being preferred when fine-tuning is done on Child Care. In two of the cases (Career and Cognitive Capabilities), the male version is preferred by more than a factor of two. Looking closely at each row of the table, we find that, after fine-tuning on Career, the model is more likely to predict the male version of a document when classifying in all other categories, except Sex and Relationships, where the male and female genders are equally represented. Similarly, when fine-tuning on Child Care, the model is significantly more likely to predict the female version of a document in all categories but Domestic Work, in which the male version is preferred by a factor of two (domestic work generally is biased towards the male gender across categories of documents). This suggests that, even on a balanced dataset in which male and female documents are represented with equal frequency, the {Sex and Relationships, Career} and Child Care categories are respectively biased towards the male and female gendered documents.

Table 4 shows the words in each category with the highest term frequency-inverse document frequency (TF-IDF) scores [11]. These results, along with the bias results shown in Table 3, illustrate that salient words like {love, emotional, friend} (for Sex and Relationships) or {job, career, home} (for Career) are incorrectly associated to the male gender, even though they are inherently gender-neutral. Likewise, {baby, parents, children} (for Child Care) are incorrectly associated to the female gender. These words lead to bias that potentially drives the model to prefer one gender over another.

Finally, we investigate the performance of our model after employing our debiasing technique. Reducing the gender gap in zero-shot conditions is challenging because we only have access to the validation accuracy on the category on which the model is fine-tuned. Tables 5 and 6 show the updated retrieval metrics after regularizing for gender difference. Our experiments showed that with strong regularization ( $\alpha_k \gg 2 \forall k$ ), we were able to bring the gender bias close to zero; however, this was at the cost of a severe degradation of retrieval performance. Thus, there is a need to balance drop in retrieval performance with reduction in gender bias. With this in mind, we fine-tuned the  $\alpha_k$ 's such that the retrieval performance did not drop by more than 2-5% while the difference between male and female bias was brought down to an average of 5.7%. Performance after fine-tuning on Cognitive Capabilities dropped the most, indicating that words in this category were especially skewed towards male.

## 6 CONCLUSION

Gender bias is a well-known, prevalent issue in many areas of NLP, including IR. Mitigating gender bias is an especially difficult problem in the context of zero-shot retrieval. In this paper, we demonstrate that employing adapter-based fine-tuning yields significantly better zero-shot retrieval performance compared to full fine-tuning of a BERT encoder. To reduce gender bias, we further introduced a debiasing regularization method, and show that we gain about 6% reduction in gender bias at the cost of 2-5% reduction in retrieval performance. Natural extensions of this work include developing methods for addressing multiple forms of bias that can be present in a single model, and in employing latent variable methods to intuit this bias when labeled data is not available.

## REFERENCES

- [1] Rishabh Bhardwaj, Navonil Majumder, and Soujanya Poria. 2021. Investigating gender bias in bert. *Cognitive Computation* 13, 4 (2021), 1008–1018.

- [2] Amin Bigdeli. 2021. Exploration and Mitigation of Stereotypical Gender Biases in Information Retrieval Systems. (2021).
- [3] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems* 29 (2016).
- [4] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- [5] Sunipa Dev and Jeff Phillips. 2019. Attenuating bias in word vectors. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 879–887.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115, 16 (2018), E3635–E3644.
- [8] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 771–787.
- [9] Klara Krieg, Emilia Parada-Cabaleiro, Gertraud Medicus, Oleg Lesota, Markus Schedl, and Navid Rekasaz. 2022. Grep-BiasIR: A Dataset for Investigating Gender Representation-Bias in Information Retrieval Results. *arXiv preprint arXiv:2201.07754* (2022).
- [10] Steve Lohr. 2018. Facial recognition is accurate, if you’re a white guy. In *Ethics of Data and Analytics*. Auerbach Publications, 143–147.
- [11] Anand Rajaraman and Jeffrey David Ullman. 2011. *Mining of massive datasets*. Cambridge University Press.
- [12] Navid Rekasaz and Markus Schedl. 2020. Do Neural Ranking Models Intensify Gender Bias?. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2065–2068.
- [13] Rashida Richardson, Jason M Schultz, and Kate Crawford. 2019. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *NYUL Rev. Online* 94 (2019), 15.
- [14] Dinghan Shen, Pengyu Cheng, Dhanasekar Sundararaman, Xinyuan Zhang, Qian Yang, Meng Tang, Asli Celikyilmaz, and Lawrence Carin. 2019. Learning compressed sentence representations for on-device text processing. *arXiv preprint arXiv:1906.08340* (2019).
- [15] Jake Silberg and James Manyika. 2019. Notes from the AI frontier: Tackling bias in AI (and in humans). *McKinsey Global Institute* (2019), 1–6.
- [16] Amit Singhal et al. 2001. Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.* 24, 4 (2001), 35–43.
- [17] Vivek Subramanian and Dhanasekar Sundararaman. 2021. How do lexical semantics affect translation? An empirical study. *arXiv preprint arXiv:2201.00075* (2021).
- [18] Dhanasekar Sundararaman, Shijing Si, Vivek Subramanian, Guoyin Wang, Devamanyu Hazarika, and Lawrence Carin. 2020. Methods for Numeracy-Preserving Word Embeddings. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 4742–4753.
- [19] Dhanasekar Sundararaman, Vivek Subramanian, Guoyin Wang, Shijing Si, Dinghan Shen, Dong Wang, and Lawrence Carin. 2019. Syntax-Infused Transformer and BERT models for Machine Translation and Natural Language Understanding. *arXiv preprint arXiv:1911.06156* (2019).
- [20] Dhanasekar Sundararaman, Vivek Subramanian, Guoyin Wang, Shijing Si, Dinghan Shen, Dong Wang, and Lawrence Carin. 2021. Syntactic Knowledge-Infused Transformer and BERT Models. In *CIKM Workshops*.
- [21] Dhanasekar Sundararaman, Vivek Subramanian, Guoyin Wang, Liyan Xu, and Lawrence Carin. 2022. Number Entity Recognition. *arXiv preprint arXiv:2205.03559* (2022).
- [22] Jieyu Zhao, Yichao Zhou, Zeyu Li, Wei Wang, and Kai-Wei Chang. 2018. Learning Gender-Neutral Word Embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 4847–4853. <https://doi.org/10.18653/v1/D18-1521>