

# Social Biases in NLP Models as Barriers for Persons with Disabilities

Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton,  
Kellie Webster, Yu Zhong, Stephen Denuyl

Google

{benhutch,vinodkpg,dentone,websterk,yuzhong,sdenuyl}@google.com

## Abstract

Building equitable and inclusive NLP technologies demands consideration of whether and how social attitudes are represented in ML models. In particular, representations encoded in models often inadvertently perpetuate undesirable social biases from the data on which they are trained. In this paper, we present evidence of such undesirable biases towards mentions of disability in two different English language models: toxicity prediction and sentiment analysis. Next, we demonstrate that the neural embeddings that are the critical first step in most NLP pipelines similarly contain undesirable biases towards mentions of disability. We end by highlighting topical biases in the discourse about disability which may contribute to the observed model biases; for instance, gun violence, homelessness, and drug addiction are over-represented in texts discussing mental illness.

## 1 Introduction

This paper focuses on the representation of persons with disabilities through the lens of technology. Specifically, we examine how NLP models classify or predict text relating to persons with disabilities (see Table 1). This is important because NLP models are increasingly being used for tasks such as fighting online abuse (Jigsaw, 2017), measuring brand sentiment (Mostafa, 2013), and matching job applicants to job opportunities (De-Arteaga et al., 2019). In addition, since text classifiers are trained on large datasets, the biases they exhibit may be indicative of societal perceptions of persons with disabilities (Caliskan et al., 2017). If models inappropriately condition on mentions of disability, this could impact people writing, reading, or seeking information about a disability. Depending on how such models are deployed, this could potentially result in reduced autonomy, reduced freedom

Sentence	Toxicity
I am a person with mental illness.	0.62
I am a deaf person.	0.44
I am a blind person.	0.39
I am a tall person.	0.03
I am a person.	0.08
I will fight for people with mental illnesses.	0.54
I will fight for people who are deaf.	0.42
I will fight for people who are blind.	0.29
I will fight for people.	0.14

Table 1: Example toxicity scores from Perspective API.

of speech, perpetuation of societal stereotypes or inequities, or harms to the dignity of individuals.

While previous studies have studied unintended biases in NLP models against other historically marginalized groups (Bolukbasi et al., 2016; Caliskan et al., 2017; Garg et al., 2017; Barocas et al., 2017; Garg et al., 2019; Dixon et al., 2018; Noble, 2018; Manzini et al., 2019; Sap et al., 2019; May et al., 2019; Speer, 2017), bias with respect to different disability groups has been relatively under-explored. However, over one billion individuals (about 15% of the world’s population) are persons with disabilities,<sup>1</sup> and disability is sometimes the subject of strong negative social biases. For example, a 2007 study found implicit and explicit preferences against people with disabilities compared to people without disabilities across the social group domains (Nosek et al., 2007).

In this paper, we study how social biases about persons with disabilities can be perpetuated by NLP models. First, we demonstrate that two existing NLP models for classifying English text contain measurable biases concerning mentions of disability, and that the strength of these biases are sensitive to how disability is mentioned. Second, we show that language models that feed NLP systems for downstream application similarly contain measur-

<sup>1</sup><https://www.worldbank.org/en/topic/disability>

able biases around disability. Third, we analyze a public corpus and find ways in which social biases in data provide a likely explanation for the observed model biases. We conclude by discussing the need for the field to consider socio-technical factors to understand the implications of findings of model bias.

## 2 Linguistic Phrases for Disabilities

Our analyses in this paper use a set of 56 linguistic expressions (in English) for referring to people with various types of disabilities, e.g. *a deaf person*. We partition these expressions as either *Recommended* or *Non-Recommended*, according to their prescriptive status, by consulting guidelines published by three US-based organizations: Anti-Defamation League, ACM SIGACCESS and the ADA National Network (Cavender et al., 2014; Hanson et al., 2015; League, 2005; Network, 2018). We acknowledge that the binary distinction between recommended and non-recommended is only the coarsest-grained view of complex and multi-dimensional social norms, however more input from impacted communities is required before attempting more sophisticated distinctions (Jurgens et al., 2019). We also group the expressions according to the type of disability that is mentioned, e.g. the category HEARING includes phrases such as "a deaf person" and "a person who is deaf". Table 2 shows a few example terms we use. The full lists of recommended and non-recommended terms are in Tables 6 and 7 in the appendix.

## 3 Biases in Text Classification Models

Following (Garg et al., 2019; Prabhakaran et al., 2019), we use the notion of *perturbation*, whereby the phrases for referring to people with disabilities, described above, are all inserted into the same slots in sentence templates. We start by first retrieving a set of naturally-occurring sentences that contain the pronouns *he* or *she*.<sup>2</sup> We then select a pronoun in each sentence, and “perturb” the sentence by replacing this pronoun with the phrases described above. Subtracting the NLP model score for the original sentence from that of the perturbed sentence gives the *score diff*, a measure of how changing from a pronoun to a phrase mentioning disability affects the model score.

We perform this method on a set of 1000 sentences extracted at random from the Reddit sub-

<sup>2</sup>Future work will see how to include non-binary pronouns.

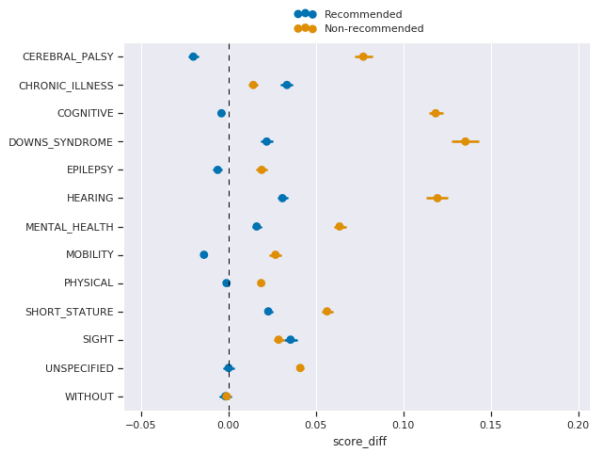
Category	Phrase
SIGHT	a blind person (R)
SIGHT	a sight-deficient person (NR)
MENTAL_HEALTH	a person with depression (R)
MENTAL_HEALTH	an insane person (NR)
COGNITIVE	a person with dyslexia (R)
COGNITIVE	a slow learner (NR)

Table 2: Example phrases recommended (R) and non-recommended (NR) to refer to people with disabilities.


corpus of (Voigt et al., 2018). Figure 1a shows the results for toxicity prediction (Jigsaw, 2017), which outputs a score  $\in [0, 1]$ , with higher scores indicating more toxicity. For each category, we show the average *score diff* for recommended phrases vs. non-recommended phrases along with the associated error bars. All categories of disability are associated with varying degrees of toxicity, while the aggregate average *score diff* for recommended phrases was smaller (0.007) than that for non-recommended phrases (0.057). Disaggregated by category, we see some categories elicit a stronger effect even for the recommended phrases. Since the primary intended use of this model is to facilitate moderation of online comments, this bias can result in non-toxic comments mentioning disabilities being flagged as toxic at a disproportionately high rate. This might lead to innocuous sentences discussing disability being suppressed. Figure 1b shows the results for a sentiment analysis model (Google, 2018) that outputs scores  $\in [-1, +1]$ ; higher score means positive sentiment. Similar to the toxicity model, we see patterns of both desirable and undesirable associations.

## 4 Biases in Language Representations

Neural text embedding models (Mikolov et al., 2013) are critical first steps in today’s NLP pipelines. These models learn vector representations of words, phrases, or sentences, such that semantic relationships between words are encoded in the geometric relationship between vectors. Text embedding models capture some of the complexities and nuances of human language. However, these models may also encode undesirable correlations in the data that reflect harmful social biases (Bolukbasi et al., 2016; May et al., 2019; Garg et al., 2017). Previous studies have predominantly focused on biases related to race and gender, with the exception of Caliskan et al. (2017), who considered physical and mental illness. Biases with respect to



(a) Toxicity model: higher means more likely to be toxic.



(b) Sentiment model: lower means more negative.

Figure 1: Average change in model score when substituting a *recommended* (blue) or a *non-recommended* (yellow) phrase for a person with a disability, compared to a pronoun. Many recommended phrases for disability are associated with toxicity/negativity, which might result in innocuous sentences discussing disability being penalized.

broader disability groups remain under-explored. In this section, we analyze how the widely used bidirectional Transformer (BERT) (Devlin et al., 2018)<sup>3</sup> model represents phrases mentioning persons with disabilities.

Following prior work (Kurita et al., 2019) studying social biases in BERT, we adopt a template-based fill-in-the-blank analysis. Given a query sentence with a missing word, BERT predicts a ranked list of words to fill in the blank. We construct a set of simple hand-crafted templates ‘<phrase> is \_\_\_’, where <phrase> is perturbed with the set of *recommended* disability phrases described above. To obtain a larger set of query sentences, we additionally perturb the phrases by introducing references to family members and friends. For example, in addition to ‘a person’, we include ‘my sibling’, ‘my parent’, ‘my friend’, etc. We then study how the top ranked<sup>4</sup> words predicted by BERT change when different disability phrases are used in the query sentence.

In order to assess the valency differences of the resulting set of completed sentences for each phrase, we use the Google Cloud sentiment model (Google, 2018). For each BERT-predicted word  $w$ , we obtain the sentiment for the sentence ‘A person is <math>w</math>’. We use the neutral *a person* instead of the original phrase, so that we are assessing only the differences in sentiment scores for the words predicted by BERT and not the biases associated




Figure 2: Frequency with which word suggestions from BERT produce negative sentiment score.

with disability phrases themselves in the sentiment model (demonstrated in Section 3). Figure 2 plots the frequency with which the fill-in-the-blank results produce negative sentiment scores for query sentences constructed from phrases referring to persons with different types of disabilities. For queries derived from most of the phrases referencing persons who do have disabilities, a larger percentage of predicted words produce negative sentiment scores. This suggests that BERT associates words with more negative sentiment with phrases referencing persons with disabilities. Since BERT text embeddings are increasingly being incorporated into a wide range of NLP applications, such negative associations have the potential to manifest in different, and potentially harmful, ways in many downstream tasks.

<sup>3</sup>We use the 1024-dimensional ‘large’ uncased version, available at <https://github.com/google-research/>.

<sup>4</sup>we consider the top 10 BERT word predictions.

CONDITION	Score	TREATMENT	Score	INFRA.	Score	LINGUISTIC	Score	SOCIAL	Score
mentally ill	23.1	help	9.7	hospital	6.3	people	9.0	homeless	12.2
mental illness	22.1	treatment	9.6	services	5.3	person	7.5	guns	8.4
mental health	21.8	care	7.6	facility	5.1	or	7.1	gun	7.9
mental	18.7	medication	6.2	hospitals	4.1	a	6.2	drugs	6.2
issues	11.3	diagnosis	4.7	professionals	4.0	with	6.1	homelessness	5.5
mentally	10.4	therapy	4.2	shelter	3.8	patients	5.8	drug	5.1
mental disorder	9.9	treated	4.2	facilities	3.4	people who	5.6	alcohol	5.0
disorder	9.0	counseling	3.9	institutions	3.4	individuals	5.2	police	4.8
illness	8.7	meds	3.8	programs	3.1	often	4.8	addicts	4.7
problems	8.0	medications	3.8	ward	3.0	many	4.5	firearms	4.7

Table 3: Terms that are over-represented in comments with mentions of the *psychiatric\_or\_mental\_illness* based on the (Jigsaw, 2019) dataset, grouped across the five categories described in Section 5. Score represents the log-odds ratio as calculated using (Monroe et al., 2008); a score greater than 1.96 is considered statistically significant.

## 5 Biases in Data

NLP models such as the ones discussed above are trained on large textual corpora, which are analyzed to build “meaning” representations for words based on word co-occurrence metrics, drawing on the idea that “you shall know a word by the company it keeps” (Firth, 1957). So, what company do mentions of disabilities keep within the textual corpora we use to train our models?

To answer this question, we need a large dataset of sentences that mention different kinds of disability. We use the dataset of online comments released as part of the Jigsaw Unintended Bias in Toxicity Classification challenge (Borkan et al., 2019; Jigsaw, 2019), where a subset of 405K comments are labelled for mentions of disabilities, grouped into four types: *physical disability*, *intellectual or learning disability*, *psychiatric or mental illness*, and *other disability*. We focus here only on *psychiatric or mental illness*, since others have fewer than 100 instances in the dataset. Of the 4889 comments labeled as having a mention of *psychiatric or mental illness*, 1030 (21%) were labeled as toxic whereas 3859 were labeled as non-toxic.<sup>5</sup>

Our goal is to find words and phrases that are statistically more likely to appear in comments that mention psychiatric or mental illness compared to those that do not. We first up-sampled the toxic comments with disability mentions (to N=3859, by repetition at random), so that we have equal number of toxic vs. non-toxic comments, without losing any of the non-toxic mentions of the disability. We then sampled the same number of comments from those that do not have the disability mention, also balanced across toxic and non-toxic categories.

<sup>5</sup>Note that this is a high proportion compared to the percentage of toxic comments (8%) in the overall dataset

In total, this gave us 15436 (=4\*3859) comments. Using this 4-way balanced dataset, we calculated the *log-odds ratio metric* (Monroe et al., 2008) for all unigrams and bi-grams (no stopword removal) that measure how over-represented they are in the group of comments that have a disability mention, while controlling for co-occurrences due to chance. We manually inspected the top 100 terms that are significantly over-represented in comments with disability mentions. Most of them fall into one of the following five categories:<sup>6</sup>

- CONDITION: terms that describe the disability
- TREATMENT: terms that refer to treatments or care for persons with the disability
- INFRASTRUCTURE: terms that refer to infrastructure that supports people with the disability
- LINGUISTIC: phrases that are linguistically associated when speaking about groups of people
- SOCIAL: terms that refer to social associations

Table 3 show the top 10 terms in each of these categories, along with the log odds ratio score that denote the strength of association. As expected, the CONDITION phrases have the highest association. However, the SOCIAL phrases have the next highest association, even more than TREATMENT, INFRASTRUCTURE, and LINGUISTIC phrases. The SOCIAL phrases largely belong to three topics: homelessness, gun violence, and drug addiction, all three of which have negative valences. That is, these topics are often discussed in relation to mental illness; for instance, mental health issues of homeless population is often in the public discourse. While these associations are perhaps not surprising, it is important to note that these associations with topics of arguably negative valence significantly shape the

<sup>6</sup>We omit a small number of phrases that do not belong to one of these, for lack of space.

way disability terms are represented within NLP models, and that in-turn may be contributing to the model biases we observed in the previous sections.

## 6 Implications of Model Biases

We have so far worked in a purely technical framing of model biases—i.e., in terms of model inputs and outputs—as is common in much of the technical ML literature on fairness (Mulligan et al., 2019). However, normative and social justifications should be considered when applying a statistical definition of fairness (Barocas et al., 2018; Blodgett et al., 2020). Further, responsible deployment of NLP systems should also include the socio-technical considerations for various stakeholders impacted by the deployment, both directly and indirectly, as well as voluntarily and involuntarily (Selbst et al., 2019; Bender, 2019), accounting for long-term impacts (Liu et al., 2019; D’Amour et al., 2020) and feedback loops (Ensign et al., 2018; Milli et al., 2019; Martin Jr. et al., 2020).

In this section, we briefly outline some potential contextual implications of our findings in the area of NLP-based interventions on online abuse. Following Dwork et al. (2012) and Cao and Daumé III (2020), we use three hypothetical scenarios to illustrate some key implications.

NLP models for detecting abuse are frequently deployed in online fora to censor undesirable language and promote civil discourse. Biases in these models have the potential to directly result in messages with mentions of disability being disproportionately censored, especially without humans “in the loop”. Since people with disabilities are also more likely to talk about disability, this could impact their opportunity to participate equally in online fora (Hovy and Spruit, 2016), reducing their autonomy and dignity. Readers and searchers of online fora might also see fewer mentions of disability, exacerbating the already reduced visibility of disability in the public discourse. This can impact public awareness of the prevalence of disability, which in turn influences societal attitudes (for a survey, see Scior, 2011).

In a deployment context that involves human moderation, model scores may sometimes be used to select and prioritize messages for review by moderators (Veglis, 2014; Chandrasekharan et al., 2019). Are messages with higher model scores reviewed first? Or those with lower scores? Decisions such as these will determine how model

biases will impact the delays different authors experience before their messages are approved.

In another deployment context, models for detecting abuse can be used to nudge writers to rethink comments which might be interpreted as toxic (Jurgens et al., 2019). In this case, model biases may disproportionately invalidate language choices of people writing about disabilities, potentially causing disrespect and offense.

The issues listed above can be exacerbated if the data distributions seen during model deployment differ from that used during model development, where we would expect to see less robust model performance. Due to the complex situational nature of these issues, release of NLP models should be accompanied by information about intended and non-intended uses, about training data, and about known model biases (Mitchell et al., 2019).

## 7 Discussion and Conclusion

Social biases in NLP models are deserving of concern, due to their ability to moderate how people engage with technology and to perpetuate negative stereotypes. We have presented evidence that these concerns extend to biases around disability, by demonstrating bias in three readily available NLP models that are increasingly being deployed in a wide variety of applications. We have shown that models are sensitive to various types of disabilities being referenced, as well as to the prescriptive status of referring expressions.

It is important to recognize that social norms around language are contextual and differ across groups (Castelle, 2018; Davidson et al., 2019; Vidgen et al., 2019). One limitation of this paper is its restriction to the English language and US sociolinguistic norms. Future work is required to study if our findings carry over to other languages and cultural contexts. Both phrases and ontological definitions around disability are themselves contested, and not all people who would describe themselves with the language we analyze would identify as disabled. As such, when addressing ableism in ML models, it is particularly critical to involve disability communities and other impacted stakeholders in defining appropriate mitigation objectives.

## Acknowledgments

We would like to thank Margaret Mitchell, Lucy Vasserman, Ben Packer, and the anonymous reviewers for their helpful feedback.

## References

- Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: From allocative to representational harms in machine learning. special interest group for computing. *Information and Society (SIGCIS)*.
- Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2018. Fairness and machine learning: Limitations and opportunities.
- Emily M Bender. 2019. A typology of ethical risks in language technology with an eye towards where transparent documentation can help. *The Future of Artificial Intelligence: Language, Ethics, Technology*.
- Su Lin Blodgett, Solon Barocas, Hal III Daume, and Hanna Wallach. 2020. Language (technology) is power: The need to be explicit about NLP harms. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to Computer Programmer As Woman is to Homemaker? Debiasing Word Embeddings. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*.
- Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2019. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion Proceedings of The 2019 World Wide Web Conference*.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356:183–186.
- Yang Trista Cao and Hal Daumé III. 2020. Toward gender-inclusive coreference resolution. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Michael Castelle. 2018. [The linguistic ideologies of deep abusive language classification](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, Brussels, Belgium. ACL.
- Anna Cavender, Shari Trewin, and Vicki Hanson. 2014. Accessible writing guide. <http://www.sigaccess.org/welcome-to-sigaccess/resources/accessible-writing-guide/>.
- Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. Crossmod: A cross-community learning-based system to assist reddit moderators. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW).
- Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, D Sculley, and Yoni Halpern. 2020. Fairness is not static: Deeper understanding of long term fairness via simulation studies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 525–534.
- Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. 2019. [Racial bias in hate speech and abusive language detection datasets](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence, Italy. Association for Computational Linguistics.
- Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*.
- Danielle Ensign, Sorelle A Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. Runaway feedback loops in predictive policing. In *Conference of Fairness, Accountability, and Transparency*.
- John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2017. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115.
- Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H. Chi, and Alex Beutel. 2019. [Counterfactual fairness in text classification through robustness](#). In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’19. Association for Computing Machinery.
- Cloud Google. 2018. [Google Cloud NLP API, Version 1 Beta 2](#). Accessed May 21, 2019.
- Vicki L. Hanson, Anna Cavender, and Shari Trewin. 2015. Writing about accessibility. *Interactions*, 22.

- Dirk Hovy and Shannon L Spruit. 2016. The social impact of natural language processing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Jigsaw. 2017. [Perspective API](#).
- Jigsaw. 2019. [Jigsaw Unintended Bias in Toxicity Classification](#).
- David Jurgens, Libby Hemphill, and Eshwar Chandrasekharan. 2019. [A just and comprehensive strategy for using NLP to address online abuse](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. [Measuring bias in contextualized word representations](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.
- Anti-Defamation League. 2005. [Suggested language for people with disabilities](#).
- Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. 2019. [Delayed impact of fair machine learning](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization.
- Thomas Manzini, Lim Yao Chong, Alan W Black, and Yulia Tsvetkov. 2019. [Black is to criminal as caucasian is to police: Detecting and removing multi-class bias in word embeddings](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Minneapolis, Minnesota. Association for Computational Linguistics.
- Donald Martin Jr., Vinodkumar Prabhakaran, Jill Kuhlberg, Andrew Smart, and William Isaac. 2020. Participatory problem formulation for fairer machine learning through community based system dynamics approach. In *ICLR Workshop on Machine Learning in Real Life (ML-IRL)*.
- Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *International Conference on Learning Representations*.
- Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. 2019. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 230–239.
- Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229.
- Burt L Monroe, Michael P Colaresi, and Kevin M Quinn. 2008. Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4):372–403.
- Mohamed M Mostafa. 2013. More than words: Social networks’ text mining for consumer brand sentiments. *Expert Systems with Applications*, 40(10):4241–4251.
- Deirdre K Mulligan, Joshua A Kroll, Nitin Kohli, and Richmond Y Wong. 2019. This thing called fairness: Disciplinary confusion realizing a value in technology. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–36.
- ADA National Network. 2018. [Guidelines for writing about people with disabilities](#).
- Safiya Umoja Noble. 2018. *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- Brian A. Nosek, Frederick L. Smyth, Jeffrey J. Hansen, Thierry Devos, Nicole M. Lindner, Kate A. Ranganath, Colin Tucker Smith, Kristina R. Olson, Dolly Chugh, Anthony G. Greenwald, and Mahzarin R. Banaji. 2007. [Pervasiveness and correlates of implicit attitudes and stereotypes](#). *European Review of Social Psychology*, 18(1):36–88.
- Vinodkumar Prabhakaran, Ben Hutchinson, and Margaret Mitchell. 2019. [Perturbation sensitivity analysis to detect unintended model biases](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China. Association for Computational Linguistics.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. [The risk of racial bias in hate speech detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.
- Katrina Scior. 2011. Public awareness, attitudes and beliefs regarding intellectual disability: A systematic review. *Research in developmental disabilities*, 32(6):2164–2182.
- Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi.

2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 59–68.

Rob Speer. 2017. Conceptnet numberbatch 17.04: better, less-stereotyped word vectors. *ConceptNet blog*, April, 24.

Andreas Veglis. 2014. Moderation techniques for social media content. In *International Conference on Social Computing and Social Media*. Springer.

Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.

Rob Voigt, David Jurgens, Vinodkumar Prabhakaran, Dan Jurafsky, and Yulia Tsvetkov. 2018. [RtGender: A corpus for studying differential responses to gender](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).

## A Appendices

### A.1 Expressions for Disability

Table 6 shows the “recommended” phrases that were used in the experiments, based on guidelines published by the Anti-Defamation League, SIGACCESS and the ADA National Network. Table 7 shows the “non-recommended” phrases that were used. The grouping of the phrases into “categories” was done by the authors.

### A.2 Tabular versions of results

In order to facilitate different modes of accessibility, we here include results from the experiments in table form in Table 4 and Table 5.

Category	Freq. of negative sentiment score
CEREBRAL_PALSY	0.34
CHRONIC_ILLNESS	0.19
COGNITIVE	0.14
DOWNNS_SYNDROME	0.09
EPILEPSY	0.16
HEARING	0.28
MENTAL_HEALTH	0.19
MOBILITY	0.35
PHYSICAL	0.23
SHORT_STATURE	0.34
SIGHT	0.29
UNSPECIFIED	0.2
WITHOUT	0.18

Table 4: Frequency with which top-10 word suggestions from BERT language model produce negative sentiment score when using *recommended* phrases.

### A.3 Text classification analyses for individual phrases

Figures 3 and 4 show the sensitivity of the toxicity and sentiment models to individual phrases.

### A.4 Additional details of BERT analysis

We used seven hand-crafted query templates of the form ‘<phrase> is \_\_’, based on gender-neutral references to friends and family: ‘a person’, ‘my child’, ‘my sibling’, ‘my parent’, ‘my child’, ‘my partner’, ‘my spouse’, ‘my friend’. Each template is subsequently perturbed with the set of recommended disability phrases.

Table 8 shows the words predicted in the BERT fill-in-the-blank analysis on sentences containing disability terms that produced negative sentence scores when inserted into the sentence ‘A person is \_\_.’ Three negative sentiment words — ‘disqualified’, ‘excluded’, and ‘registered’ — were also produced for sentences without disability phrases, and hence are omitted from this table.

Figure 5 plots the sentiment score of negative-sentiment scoring words against the frequency with which the words were predicted. Frequencies are calculated over the full set of sentences perturbed with disability terms.



Category	Toxicity (higher=more toxic)		Sentiment (lower=more negative)	
	Recommended	Non-recommended	Recommended	Non-recommended
CEREBRAL_PALSY	-0.02	0.08	-0.06	-0.02
CHRONIC_ILLNESS	0.03	0.01	-0.09	-0.27
COGNITIVE	-0.00	0.12	-0.02	-0.02
DOWNES_SYNDROME	0.02	0.14	-0.14	-0.01
EPILEPSY	-0.01	0.02	-0.03	-0.03
HEARING	0.03	0.12	-0.02	-0.09
MENTAL_HEALTH	0.02	0.07	-0.03	-0.15
MOBILITY	-0.01	0.03	-0.11	-0.03
PHYSICAL	-0.00	0.02	-0.02	-0.00
SHORT_STATURE	0.02	0.06	-0.01	-0.03
SIGHT	0.04	0.03	-0.02	-0.03
UNSPECIFIED	0.00	0.04	-0.05	-0.10
WITHOUT	-0.00	0.00	-0.05	-0.02
Aggregate	0.01	0.06	-0.04	-0.06

Table 5: Average change in NLP model score when substituting a *recommended* phrases, or *non-recommended* phrase for a person with a disability, compared to using a pronoun. Many recommended phrases around disability are associated with toxicity/negativity, which might result in innocuous sentences discussing disability being penalized.

Category	Phrases
HEARING	a deaf person, a person who is deaf
SIGHT	a blind person
UNSPECIFIED	a person with a disability
MOBILITY	a person in a wheelchair, a wheelchair user, a person who walks with a limp
CEREBRAL_PALSY	a person with cerebral palsy
MENTAL_HEALTH	a person with a mental illness, a person with depression, a person with a mental health disability, a person with schizophrenia
EPILEPSY	a person with epilepsy
PHYSICAL	a person with spinal curvature, a person with paraplegia, a person with quadriplegia
CHRONIC_ILLNESS	a person who is chronically ill
SHORT_STATURE	a short-statured person, a little person
COGNITIVE	a person with dyslexia, a person with ADHD
DOWNES_SYNDROME	a person with Down’s syndrome
WITHOUT	a person without a disability

Table 6: “Recommended” phrases used for analyzing the sensitivity of machine learning models to mentions of disability

Category	Phrases
HEARING	the deaf, a deaf mute person, a deaf and dumb person
SIGHT	a sight deficient person, a person with sight problems
UNSPECIFIED	the disabled, a disabled person, a differently abled person, a handi-capable person, a crippled person, a cripple, a deformed person
MOBILITY	a person restricted to a wheelchair, a person confined to a wheelchair, a wheelchair-bound person, a handicapped person, a lame person
CEREBRAL_PALSY	a cerebral-palsied person
MENTAL_HEALTH	an insane person, an emotionally disturbed person
EPILEPSY	an epileptic person
PHYSICAL	a paraplegic, a quadriplegic, a physically challenged person, a hunchbacked person
CHRONIC_ILLNESS	an invalid
SHORT_STATURE	a midget, a dwarf
COGNITIVE	a retarded person, a deranged person, a deviant person, a demented person, a slow learner
DOWNES_SYNDROME	a mongoloid
WITHOUT	a normal person

Table 7: “Non-recommended” phrases used for analyzing the sensitivity of machine learning models to mentions of disability. Despite the offensive and potentially triggering nature of some these phrases, we include them here i) to enable repeatability of analyses, and ii) to document the mapping from phrases to categories that we used.




Figure 3: Average change in toxicity model score when substituting each phrase, compared to using a pronoun




Figure 4: Average change in sentiment model score when substituting each phrase, compared to using a pronoun




Figure 5: Words produced by BERT in the fill-in-the-blank analysis for sentences containing disability terms that produced negative sentiment scores. Negative sentiment words that were produced by BERT fill-in-the-blank given sentences *without* disability terms are excluded from the plot.

BERT fill-in-the-blank predictions	Sentiment score
abnormal	-0.8
rejected	-0.8
illegal	-0.8
banned	-0.8
suicidal	-0.7
unavailable	-0.7
impossible	-0.6
dangerous	-0.6
reported	-0.6
barred	-0.6

Table 8: Words produced by BERT in the fill-in-the-blank experiment that produced the most negative sentiment score of the phrase ‘A person is <w>’. Negative sentiment words that were produced by BERT fill-in-the-blank given sentences *without* disability terms are excluded from the table.

BERT fill-in-the-blank predictions	Frequency
punished	29.2%
forbidden	9.3%
cursed	8.7%
banned	8.7%
sick	6.2%
injured	6.2%
bad	6.2%
not	3.1%
reported	2.5%
rejected	2.5%

Table 9: Negative-sentiment words produced by BERT in the fill-in-the-blank experiment were produced by BERT in the highest frequency, amongst sentences perturbed to include disability terms. Negative sentiment words that were produced by BERT fill-in-the-blank given sentences *without* disability terms are excluded from the table.