

# What’s in the Box?

## A Preliminary Analysis of Undesirable Content in the Common Crawl Corpus

**Alexandra (Sasha) Luccioni**

Université de Montréal &

Mila Québec AI Institute

sasha.luccioni@mila.quebec

**Joseph D. Viviano**

Mila Québec AI Institute

joseph@viviano.ca

### Abstract

Whereas much of the success of the current generation of neural language models has been driven by increasingly large training corpora, relatively little research has been dedicated to analyzing these massive sources of textual data. In this exploratory analysis, we delve deeper into the Common Crawl, a colossal web corpus that is extensively used for training language models. We find that it contains a significant amount of undesirable content, including hate speech and sexually explicit content, even after filtering procedures. We discuss the potential impacts of this content on language models and conclude with future research directions and a more mindful approach to corpus collection and analysis.

### 1 Introduction

In recent years, much of the progress in Natural Language Processing (NLP) research has been largely driven by Transformer-based language models, which have pushed forward the state-of-the-art in tasks such as question answering (Rajpurkar et al., 2018) and natural language inference (Bowman et al., 2015). However, these increasingly complex models also require increasingly large amounts of data to train them, which is often a combination of curated, high-quality datasets such as encyclopedic articles and books and non-curated content from the Web (Radford et al., 2018, 2019). This second category of large, non-curated dataset is becoming increasingly popular as they are required to train large language models.

The current largest dataset used for training neural language models, the Common Crawl, is a non-curated corpus consisting of multilingual snapshots of the web. New versions of the Common Crawl are released monthly, with each version containing 200 to 300 TB of textual content scraped via automatic web crawling. This dwarfs other commonly used corpora such as English-language

Wikipedia, which adds up to roughly 5.6 TB of data, and the BookCorpus, which only represents around 6 GB (Zhu et al., 2015). The Common Crawl has been used to train many of the recent neural language models in recent years, including the GPT model series (Radford et al., 2018; Brown et al., 2020), BERT (Devlin et al., 2018) and FastText (Grave et al., 2018) and, given its size, often represents the majority of data used to train these architectures.

In the current article, we present an initial analysis of the Common Crawl, highlighting the presence of several types of explicit and abusive content even after filtering. We discuss our findings and, given the potential downstream impact of this content on language models, we discuss the importance of ensuring that the corpora we use for training language models are extracted more mindfully and with more emphasis on their quality and propose avenues of research to achieve this goal.

### 2 Related Work

In recent years, a growing body of research in NLP has unearthed biases in common language models (Bolukbasi et al., 2016; Sheng et al., 2019; Zhao et al., 2019; Bordia and Bowman, 2019; Hutchinson et al., 2020). This work has raised important questions regarding the impact of these embedded biases on downstream decision-making, given the increasing usage of these models in various applications. Consequently, much work has also been dedicated to creating standardized diagnostic tests to detect these biases (Caliskan et al., 2017; May et al., 2019; Nadeem et al., 2020; Sweeney and Najafian, 2019) and to remove them (Bolukbasi et al., 2016; Zhao et al., 2018; Manzini et al., 2019), although the extent to which this is possible is still under debate (Gonen and Goldberg, 2019). In fact, research has found that “*The biases found in Internet-scale language models like GPT-2 are representative of the data on which the model was trained*” (So-

laiman et al., 2019), which can be directly linked to the presence of hate speech on the Internet (Abid et al., 2021).

However, given the importance of this research, comparatively little attention has been dedicated to analyzing the corpora used to train language models. This is understandable because frequently used datasets such as the Common Crawl contain truly massive amounts of data, making it challenging to mine it for meaningful insights. In fact, a recent survey on automatic web page classification has deemed the task difficult not only due to the complexity and heterogeneity of web content, but also due its the high computational cost, suggesting that machine learning (ML) approaches have much to contribute to it (Hashemi, 2020). While certain notable endeavors have indeed analyzed specific aspects of corpora such as the Common Crawl (Kolias et al., 2014; Caswell et al., 2021) and Wikipedia (Hube, 2017), they have only scratched the surface of what these bodies of text contain. For instance, recent work has found that the Common Crawl contained over 300,000 documents from unreliable news sites and banned subReddit pages containing hate speech and racism (Gehman et al., 2020), while complementary research has shown that individual training examples can be extracted by querying language models (Carlini et al., 2020), together illustrating that the presence of questionable content is a significant issue for statistical language models. In the current work, we endeavor to understand the content and quality of the Common Crawl as a first step towards establishing more consistent approaches to filtering and refining it.

### 3 Analyzing the Common Crawl

Given its size, both downloading and analyzing the Common Crawl are time-consuming and costly endeavors. The [most recent version](#) of the Common Crawl, dating from November/December 2020, has 2.6 billion web pages in raw text format, saved in ‘shards’ each containing of tens of thousands of pages. Given our hardware constraints, we chose to focus on a subset of the corpus, randomly sampling 1% of the files it contains, which after filtering by language amounts to roughly 115 GB of textual content or 5,835,339 web pages in total, which we analyzed in terms of hate speech, adult content, and efficacy of perplexity-based filtering <sup>1</sup>. In this work,

<sup>1</sup>All code used in these analysis are publicly available: <https://github.com/josephdiviviano/whatsinthebox>

we focus on detecting sexually-explicit and hate speech, since they represent common examples of “undesirable” content that can be generally seen as inappropriate for a language model to generate in most situations. We acknowledge that desirable model behaviour is application specific, and believe our findings can extend to any other “undesirable” topic that might be present in available language corpora. We present our results in the sections below.

#### 3.1 Detecting Hate Speech

The existence of hate speech on the internet has been described as “an important societal problem of our time”, with “profound and lasting” psychological effects on its victims (Mishra et al., 2019). As such, a substantial amount of NLP research dedicated to automating hate speech detection, with several datasets and approaches being proposed in recent years (Schmidt and Wiegand, 2017; Mishra et al., 2019; Vidgen and Derczynski, 2020; Kiritchenko and Mohammad, 2018). Most of this research is carried out on data extracted from social media sources such as Twitter (Founta et al., 2018; Basile et al., 2019; Waseem and Hovy, 2016) and Reddit (Tadesse et al., 2019; Farrell et al., 2019), with both ML-based (Badjatiya et al., 2017) and count-based approaches (Davidson et al., 2017) achieving comparable results (Fortuna and Nunes, 2018). In order to estimate the quantity of hate speech in the Common Crawl, we endeavored to compare 3 approaches: DELIMIT, a recent BERT-based model trained on social media data (Aluru et al., 2020), Hate Sonar, a Logistic Regression approach trained on data from Web fora and Twitter (Davidson et al., 2017) and a n-gram-based approach using a list of n-grams extracted from [Hate Base](#). We present samples of text flagged by all of these approaches in Table 1, below.

We found that the three approaches compared suggest similar proportions of websites containing hate speech : 5.24% of websites from our sample were flagged by DELIMIT, 4.02% by HateSonar, and 6.38% by the n-gram approach <sup>2</sup>. Qualitative analysis of a sample of sites flagged by each approach showed that while n-grams picked up on racial slurs, HateSonar also detected debates about racial supremacy and racially-charged conspiracy theories. Many of the sites that DELIMIT

<sup>2</sup>We are conscious of the high false positive rate of n-gram approaches and therefore only consider sites to be flagged if they contain 3 or more n-grams from the list.

Approach	Text
<b>HateSonar</b>	Their US/Euro plan put in your face: demonic jews hate white goyim! Such sick and twisted people, white people are.
<b>Delimit</b>	they are only stupid arab from wp-ar haha Yeah, dumb ass n*gger †
<b>N-gram</b>	nude attention whore asian bastards In America all male look like this homo

Table 1: Examples of hate speech found by the approaches tested. Examples with † have been censored by the authors.

flagged were adult content with mentions of violent acts towards specific ethnic groups, illustrating the fine line between sexual violence and hate speech, which we elaborate further in the following subsection. Generally speaking, the presence of even a small fraction of websites that incite hate in training corpora is worrisome since it can result in models that replicate this kind of discourse when prompted (Wolf et al., 2017; Carlini et al., 2020).

### 3.2 Sexually Explicit Content

Compared to hate speech, the detection of sexually explicit content has received less attention from the NLP community, with existing ML approaches focusing mainly on the detection of explicit images (Wehrmann et al., 2018; Rowley et al., 2006) and URLs (Matic et al., 2020), whereas n-gram-based approaches remain predominantly used in practice by web providers (Hammami et al., 2003; Polpinij et al., 2006; Ho and Watters, 2004). In our analysis, we used a list of n-grams extracted from adult websites in order to establish the percentage of websites from our sample that contained sexually explicit content; however, we found no available statistical or ML-based approach that we could use to compare our count-based approach with. The n-gram approach detected that 2.36% of the web pages that we analyzed contained at least one of the words from our list, with 1.36% containing 3 or more and 0.73% containing 10 or more (see Table 3 for results). We show a sample of the URLs flagged by our approach in Table 2, below.

While a few percent of sexually explicit content may not seem like much, the type of language and content contained on adult websites can have harmful repercussions. For instance, the prevalence of sexual violence towards women, especially towards women of color, on adult websites (Foubert et al.,

Page URL ( <a href="http://removed">http:// removed</a> )
adultmovietop100.com/ erohon.me/ celebrityfan.net/ queantube.com/ adelaide-femaleescorts.webcam

Table 2: Sample of URLs of adult content websites identified by the n-gram approach. Protocol removed to prevent URL generation.

2019; Shim et al., 2015; Fritz et al., 2020) may contribute to further dissemination and amplification of these biases in downstream models. As modern language models have no way to evaluate generation appropriateness, models trained with even a small proportion of these undesirable inputs cannot be guaranteed to avoid generating outputs with similar biases if presented with a specific context or prompt. This is a risk that is important to mitigate in applications, where the general-purpose language models can end up being used in applications used by sensitive groups in professional contexts or minors, such as chatbots and toys.

### 3.3 Filtering by Perplexity Score

While the analyses described above were carried out on unfiltered web pages from the Common Crawl, the training pipeline of many large-scale NLP models involves some type of filtering and cleaning, from excluding low-quality content (Grave et al., 2018) to fuzzy deduplication (Brown et al., 2020). One such popular filtering approach is based on training a language model on a target, high-quality domain such as Wikipedia, and using it to calculate the perplexity score of web pages using this model (Wenzek et al., 2020). To test the efficacy of this scoring procedure, we calculated the perplexity score of each web page from our sample of the Common Crawl and used it to separate pages into 3 equal buckets (high, middle and low-quality) based on their perplexity. We compare the percentages of hate speech and sexually explicit content for the entire sample, as well as the high- and low-quality documents, in Table 3.

While filtering by perplexity does seem to filter out many websites containing sexual content, it does not detect much of the hate speech that is flagged by the count-based or statistical methods. In fact, perplexity scores had low correlations with all detection methods tested (Figure 1). This supports the methodology of Wenzek et al. (2020),

	Entire Sample	High Quality	Low Quality
<b>1+ sexual n-grams</b>	2.36%	1.81%	3.97%
<b>3+ sexual n-grams</b>	1.36%	0.42%	3.11%
<b>10+ sexual n-grams</b>	0.73%	0.08%	1.98%
<b>1+ hate n-grams</b>	17.78%	18.95%	17.19%
<b>3+hate n-grams</b>	6.38%	6.19%	8.26%
<b>10+ hate n-grams</b>	1.16%	1.17%	1.70%
<b>Hate speech (Sonar)</b>	4.02%	3.47%	5.09%
<b>Hate speech (Delimit)</b>	5.24%	5.77%	5.66%

Table 3: Comparison of hate speech and sexual content detected in the entire corpus, as well as high- and low-quality sites.

who noted that while “*perplexity was a relative good proxy for quality*”, also argued that some of the lower-quality texts could still be useful for specific applications, and therefore did not use it to exclude documents from the training set of their language model. While we are exploring ways of modifying the original approach in order to be more discerning, we believe that there more nuanced metrics that can be used for estimating and filtering documents based on text, potentially coupling embedding-based approaches with statistical ones.

### 3.4 Behaviour of Different Detection Methods

The approaches that we compared in the current study are different in the features that they use and techniques employed for detecting particular types of content. HateSonar employs classical NLP techniques for hate speech detection, constructing features from Penn Part-of-Speech N-grams with TF-IDF weighting based on a hand-crafted hate speech dataset, training simple classifier ensembles using Support Vector Machines, random forests, naive Bayes, and linear models. Delimit, on the other hand, is a BERT-based model trained on Twitter and Reddit posts, not relying on any handcrafted features. Our simple n-gram approach unsurpris-

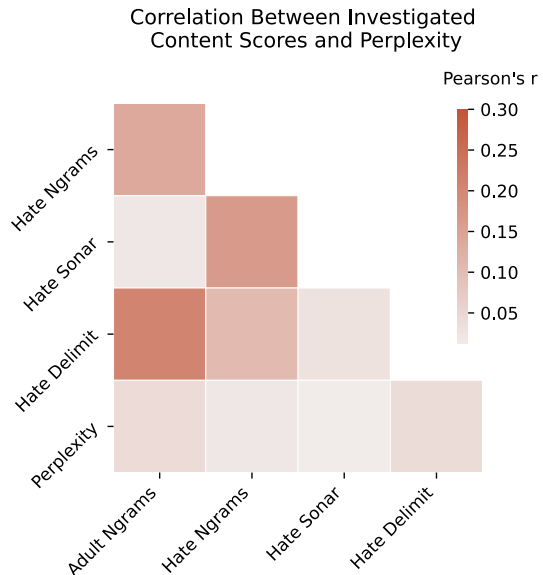


Figure 1: Correlation coefficients (Pearson’s  $r$ ) calculated between all content metrics investigated and perplexity, a commonly-used text quality metric.

ingly was more in agreement with HateSonar than Delimit, given that both rely on count-based features. The fact that all methods identified different instances of clear hate speech implies that we are far from a general purpose dataset-filtering approach. These results also imply that deep learning models learn very different features to classify hate speech than other methods, and given their sensitivity to the specific composition of the dataset used to train them (as exposed by the propensity of large models to memorize training examples (Carlini et al., 2020)), the presence of undesirable content in the corpora used to train them should be taken seriously.

## 4 Discussion

### 4.1 Summary of Results

We recognize that the exploratory work presented above is only the tip of the iceberg in terms of the analyses that can be done on the massive web corpora that are feeding our language models. However, analyzing the Common Crawl would require computational resources far in excess of what is available to most research institutions. We therefore hope that this initial analysis will inspire our fellow researchers to continue to dig deeper into this topic, and to propose more scalable, thorough, and nuanced approaches for analyzing the massive corpora used to train language models. We also recognize this analysis would have been more comprehensive on a small curated dataset, but given the

amount of data needed to train modern language models, we believe the community needs to move beyond analysis techniques only compatible with small-data, toward something that will scale to the datasets used to train these large models.

Also, while we have currently adopted a purely descriptive approach, we feel that it is worth discussing and debating the consequences of our analysis, and those of our peers, within the NLP community. While it can be argued that the Common Crawl corpus is an accurate portrayal of the discourse of modern society – which includes sexual content, hate speech, and racial and gender biases – we believe that it is up for debate whether this discourse is the one that we, as a community, want to use to train the models that translate our texts, influence our search results and answer our questions. Notably, the Common Crawl over-represents those populations that are avid users of the internet: younger, English-speaking individuals from developed countries, who are those who have the most access to the internet globally (World Bank, 2018). Furthermore, internet communities supported by anonymity and particular norms can amplify toxic discourse that would not be found in mainstream corpora (Massanari, 2017) often exacerbated by the well-documented ‘online disinhibition’ phenomenon where users find themselves more likely to engage in anti-social behaviours due to the lack of immediate social feedback (Wachs et al., 2019; Mathew et al., 2019; de Lima et al., 2021). This can further perpetuate the lack of diverse, representative language models that can adequately mirror society beyond the boundaries of internet communities.

## 4.2 Future Work

Given the general superior performance of large language models on common benchmarks, and that they require ever larger datasets to train them, we believe it is important that for the ML community to carry out a more extensive analysis of: 1) the impact of undesirable content in the datasets used to train these models on downstream performance; 2) the effect of properly filtering these examples out of the dataset *before* model training, and 3) approaches for regularizing model outputs to be acceptable regardless of the data used to train the model. All three directions require a better understanding of the contents of the datasets, which we believe requires new tools that are scalable to the

Common Crawl (or similarly large and diverse corpora) to identify such examples. Models trained to detect undesirable examples, like the ones used in this paper, need to be improved such that they can reliably generalize to the Common Crawl, which constitutes a significant undertaking. Additionally, future work could explore the utility of controlling model generation using labelled “undesirable” examples (Zhang et al., 2020; Engel et al., 2017), or human-in-the-loop learning methods (Wang et al., 2021) for fine-tuning a language model trained using undesirable examples. It will also be important to evaluate whether curation is sufficient: it remains possible that a model could create an undesirable generation from multiple distinct innocuous examples (Bender et al., 2021; Gehman et al., 2020). It is also worth considering that for some applications, task-focused models with curated training examples may perform better than large models trained on unfiltered corpora, so that their behaviour can be more reliably guaranteed: these are all interesting avenues for future work.

Finally, while larger corpora generally result in better models (Kaplan et al., 2020; Sun et al., 2017), data quality and corpora content also plays a major role in the caliber and appropriateness of these models for the various downstream applications (Florez, 2019; Abid et al., 2021; Bhardwaj et al., 2021). To produce high quality and safe neural language models will likely require the community to adopt more mindful data collection practices (Gehman et al., 2020; Bender and Friedman, 2018; Gebru et al., 2018; Jo and Gebru, 2020; Paullada et al., 2020; Bender et al., 2021), establish standardized filtering pipelines for corpora (Roziewski and Stokowiec, 2016; Ortiz Suarez et al., 2019; Wenzek et al., 2020), and develop methods for evaluating the bias in trained models (Schick et al., 2021). We recognize that this is not a straightforward task with a one-size-fits all solution, but we propose that as much attention should be dedicated to the corpora used for training language models as to the models themselves, and that corpora transparency is a prerequisite for language model accountability.

## References

- Abid, A., Farooqi, M., and Zou, J. (2021). Persistent Anti-Muslim Bias in Large Language Models. *arXiv preprint arXiv:2101.05783*.
- Aluru, S. S., Mathew, B., Saha, P., and Mukherjee, A. (2020). Deep learning models for multilingual hate speech detection. *arXiv preprint arXiv:2004.06465*.
- Badjatiya, P., Gupta, S., Gupta, M., and Varma, V. (2017). Deep learning for hate speech detection in tweets. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 759–760.
- Basile, V., Bosco, C., Fersini, E., Debora, N., Patti, V., Pardo, F. M. R., Rosso, P., Sanguinetti, M., et al. (2019). Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter. In *13th International Workshop on Semantic Evaluation*, pages 54–63. Association for Computational Linguistics.
- Bender, E., Gebru, T., McMillan-Major, A., et al. (2021). On the dangers of stochastic parrots: Can language models be too big. *Proceedings of FAccT*.
- Bender, E. M. and Friedman, B. (2018). Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Bhardwaj, R., Majumder, N., and Poria, S. (2021). Investigating gender bias in BERT. *Cognitive Computation*, pages 1–11.
- 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 *Advances in Neural Information Processing Systems*, pages 4349–4357.
- Bordia, S. and Bowman, S. R. (2019). Identifying and reducing gender bias in word-level language models. *arXiv preprint arXiv:1904.03035*.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015). A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Caliskan, A., Bryson, J. J., and Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. (2020). Extracting training data from large language models. *arXiv preprint arXiv:2012.07805*.
- Caswell, I., Kreutzer, J., Wang, L., Wahab, A., van Esch, D., Ulzii-Orshikh, N., Tapo, A., Subramani, N., Sokolov, A., Sikasote, C., et al. (2021). Quality at a glance: An audit of web-crawled multilingual datasets. *arXiv preprint arXiv:2103.12028*.
- Davidson, T., Warmsley, D., Macy, M., and Weber, I. (2017). Automated hate speech detection and the problem of offensive language. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 11.
- de Lima, L. H. C., Reis, J., Melo, P., Murai, F., and Benevenuto, F. (2021). Characterizing (un) moderated textual data in social systems. *arXiv preprint arXiv:2101.00963*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Engel, J., Hoffman, M., and Roberts, A. (2017). Latent constraints: Learning to generate conditionally from unconditional generative models. *arXiv preprint arXiv:1711.05772*.
- Farrell, T., Fernandez, M., Novotny, J., and Alani, H. (2019). Exploring misogyny across the manosphere in reddit. In *Proceedings of the 10th ACM Conference on Web Science*, pages 87–96.
- Florez, O. U. (2019). On the unintended social bias of training language generation models with data from local media. *arXiv preprint arXiv:1911.00461*.
- Fortuna, P. and Nunes, S. (2018). A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- Foubert, J. D., Blanchard, W., Houston, M., and Williams, R. R. (2019). Pornography and sexual violence. In *Handbook of Sexual Assault and Sexual Assault Prevention*, pages 109–127. Springer.
- Founta, A., Djouvas, C., Chatzakou, D., Leontiadis, I., Blackburn, J., Stringhini, G., Vakali, A., Sirivianos, M., and Kourtellis, N. (2018). Large scale crowdsourcing and characterization of Twitter abusive behavior. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 12.
- Fritz, N., Malic, V., Paul, B., and Zhou, Y. (2020). Worse than objects: The depiction of black women and men and their sexual relationship in pornography. *Gender Issues*, pages 1–21.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., and Crawford, K. (2018). Datasheets for datasets. *arXiv preprint arXiv:1803.09010*.

- Gehman, S., Gururangan, S., Sap, M., Choi, Y., and Smith, N. A. (2020). Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462*.
- Gonen, H. and Goldberg, Y. (2019). Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. *arXiv preprint arXiv:1903.03862*.
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., and Mikolov, T. (2018). Learning word vectors for 157 languages. *arXiv preprint arXiv:1802.06893*.
- Hammami, M., Chahir, Y., and Chen, L. (2003). Web-guard: Web based adult content detection and filtering system. In *Proceedings IEEE/WIC International Conference on Web Intelligence (WI 2003)*, pages 574–578. IEEE.
- Hashemi, M. (2020). Web page classification: a survey of perspectives, gaps, and future directions. *Multi-media Tools and Applications*, pages 1–25.
- Ho, W. H. and Watters, P. A. (2004). Statistical and structural approaches to filtering internet pornography. In *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583)*, volume 5, pages 4792–4798. IEEE.
- Hube, C. (2017). Bias in Wikipedia. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pages 717–721.
- Hutchinson, B., Prabhakaran, V., Denton, E., Webster, K., Zhong, Y., and Denuyl, S. (2020). Social biases in NLP models as barriers for persons with disabilities. *arXiv preprint arXiv:2005.00813*.
- Jo, E. S. and Gebru, T. (2020). Lessons from archives: strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 306–316.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Kiritchenko, S. and Mohammad, S. M. (2018). Examining gender and race bias in two hundred sentiment analysis systems. *arXiv preprint arXiv:1805.04508*.
- Kolias, V., Anagnostopoulos, I., and Kayafas, E. (2014). Exploratory analysis of a terabyte scale web corpus. *arXiv preprint arXiv:1409.5443*.
- Manzini, T., Lim, Y. C., Tsvetkov, Y., and Black, A. W. (2019). Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in word embeddings. *arXiv preprint arXiv:1904.04047*.
- Massanari, A. (2017). # gamergate and the fapping: How reddit’s algorithm, governance, and culture support toxic technocultures. *New media & society*, 19(3):329–346.
- Mathew, B., Illendula, A., Saha, P., Sarkar, S., Goyal, P., and Mukherjee, A. (2019). Temporal effects of unmoderated hate speech in gab. *arXiv preprint arXiv:1909.10966*.
- Matic, S., Iordanou, C., Smaragdakis, G., and Laoutaris, N. (2020). Identifying sensitive URLs at web-scale. In *Proceedings of the ACM Internet Measurement Conference*, pages 619–633.
- May, C., Wang, A., Bordia, S., Bowman, S. R., and Rudinger, R. (2019). On measuring social biases in sentence encoders. *arXiv preprint arXiv:1903.10561*.
- Mishra, P., Yannakoudakis, H., and Shutova, E. (2019). Tackling online abuse: A survey of automated abuse detection methods. *arXiv preprint arXiv:1908.06024*.
- Nadeem, M., Bethke, A., and Reddy, S. (2020). Stereoset: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456*.
- Ortiz Suarez, P. J., Sagot, B., and Romary, L. (2019). Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures. *Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019*. Cardiff, 22nd July 2019, pages 9 – 16, Mannheim. Leibniz-Institut für Deutsche Sprache.
- Paullada, A., Raji, I. D., Bender, E. M., Denton, E., and Hanna, A. (2020). Data and its (dis) contents: A survey of dataset development and use in machine learning research. *arXiv preprint arXiv:2012.05345*.
- Polpinij, J., Chotthanom, A., Sibunruang, C., Chamchong, R., and Puangpronpitag, S. (2006). Content-based text classifiers for pornographic web filtering. In *2006 IEEE International Conference on Systems, Man and Cybernetics*, volume 2, pages 1481–1485. IEEE.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Rajpurkar, P., Jia, R., and Liang, P. (2018). Know what you don’t know: Unanswerable questions for SQuAD. *arXiv preprint arXiv:1806.03822*.
- Rowley, H. A., Jing, Y., and Baluja, S. (2006). Large scale image-based adult-content filtering. *Google Research Paper*.
- Roziewski, S. and Stokowiec, W. (2016). Language-crawl: A generic tool for building language models upon common-crawl. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2789–2793.

- Schick, T., Udupa, S., and Schütze, H. (2021). Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp. *arXiv preprint arXiv:2103.00453*.
- Schmidt, A. and Wiegand, M. (2017). A survey on hate speech detection using natural language processing. In *Proceedings of the Fifth International workshop on natural language processing for social media*, pages 1–10.
- Sheng, E., Chang, K.-W., Natarajan, P., and Peng, N. (2019). The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326*.
- Shim, J. W., Kwon, M., and Cheng, H.-I. (2015). Analysis of representation of sexuality on women’s and men’s pornographic websites. *Social Behavior and Personality: an international journal*, 43(1):53–62.
- Solaiman, I., Brundage, M., Clark, J., Askell, A., Herbert-Voss, A., Wu, J., Radford, A., Krueger, G., Kim, J. W., Kreps, S., et al. (2019). Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.
- Sun, C., Shrivastava, A., Singh, S., and Gupta, A. (2017). Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852.
- Sweeney, C. and Najafian, M. (2019). A transparent framework for evaluating unintended demographic bias in word embeddings. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1662–1667.
- Tadesse, M. M., Lin, H., Xu, B., and Yang, L. (2019). Detection of depression-related posts in reddit social media forum. *IEEE Access*, 7:44883–44893.
- Vidgen, B. and Derczynski, L. (2020). Directions in abusive language training data: Garbage in, garbage out. *arXiv preprint arXiv:2004.01670*.
- Wachs, S., Wright, M. F., and Vazsonyi, A. T. (2019). Understanding the overlap between cyberbullying and cyberhate perpetration: Moderating effects of toxic online disinhibition. *Criminal Behaviour and Mental Health*, 29(3):179–188.
- Wang, Z. J., Choi, D., Xu, S., and Yang, D. (2021). Putting humans in the natural language processing loop: A survey. *arXiv preprint arXiv:2103.04044*.
- Waseem, Z. and Hovy, D. (2016). Hateful symbols or hateful people? predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL student research workshop*, pages 88–93.
- Wehrmann, J., Simões, G. S., Barros, R. C., and Cavalcante, V. F. (2018). Adult content detection in videos with convolutional and recurrent neural networks. *Neurocomputing*, 272:432–438.
- Wenzek, G., Lachaux, M.-A., Conneau, A., Chaudhary, V., Guzmán, F., Joulin, A., and Grave, É. (2020). CCNet: Extracting high quality monolingual datasets from web crawl data. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4003–4012.
- Wolf, M. J., Miller, K. W., and Grodzinsky, F. S. (2017). Why we should have seen that coming: comments on microsoft’s tay “experiment,” and wider implications. *The ORBIT Journal*, 1(2):1–12.
- World Bank (2018). Individuals using the Internet. <https://data.worldbank.org/indicator/IT.NET.USER.ZS?end=2017&locations=US&start=2015>. Accessed: 2021-01-10.
- Zhang, Y., Wang, G., Li, C., Gan, Z., Brockett, C., and Dolan, B. (2020). Pointer: Constrained text generation via insertion-based generative pre-training. *arXiv preprint arXiv:2005.00558*.
- Zhao, J., Wang, T., Yatskar, M., Cotterell, R., Ordonez, V., and Chang, K.-W. (2019). Gender bias in contextualized word embeddings. *arXiv preprint arXiv:1904.03310*.
- Zhao, J., Zhou, Y., Li, Z., Wang, W., and Chang, K.-W. (2018). Learning gender-neutral word embeddings. *arXiv preprint arXiv:1809.01496*.
- Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27.