

AN ANALYSIS OF THE EFFECTS OF DECODING ALGORITHMS ON FAIRNESS IN OPEN-ENDED LANGUAGE GENERATION

Jwala Dhamala^{*1}, Varun Kumar^{*1}, Rahul Gupta¹, Kai-Wei Chang^{1,2}, Aram Galstyan^{1,3}

¹Amazon, Alexa AI, USA

²University of California, Los Angeles

³Information Sciences Institute, University of Southern California

ABSTRACT

Several prior works have shown that language models (LMs) can generate text containing harmful social biases and stereotypes. While decoding algorithms play a central role in determining properties of LM generated text, their impact on the fairness of the generations has not been studied. We present a systematic analysis of the impact of decoding algorithms on LM fairness, and analyze the trade-off between fairness, diversity and quality. Our experiments with top- p , top- k and temperature decoding algorithms, in open-ended language generation, show that fairness across demographic groups changes significantly with change in decoding algorithm's hyper-parameters. Notably, decoding algorithms that output more diverse text also output more texts with negative sentiment and regard. We present several findings and provide recommendations on standardized reporting of decoding details in fairness evaluations and optimization of decoding algorithms for fairness alongside quality and diversity.

Index Terms—Language Models, Fairness, Bias, Natural Language Generation, Decoding algorithms

1. INTRODUCTION

Generating coherent and fluent text that is indistinguishable from human written text is one of the grand goals in natural language generation (NLG). The advent of language models (LMs) trained on massive scale data such as GPT-2 [1] and GPT-3 [2] have taken us closer towards achieving this goal.

Decoding algorithms use the probability distribution from LM to control how it outputs a sequence of words. While automatic evaluation of this machine generated text remains unresolved, the NLP community has primarily focused on: (1) *quality*, and (2) *diversity* when evaluating or developing these algorithms. Early studies focused on improving the quality of generations [3, 4]. Most recent ones focus on improving both quality and diversity [5, 6]. For example, to balance the quality-diversity trade-off, [7] propose temperature sweep, [8] study top- k decoding, and [9] develop nucleus sampling as an improvement over top- k .

In parallel, various works have shown that text generated by LMs contain harmful biases, such as stereotypes [10, 5], negative sentiments, and toxicity [11, 12, 13] towards historically marginalized demographic groups. These biases, when propagated to downstream tasks, can result in disparate treatment and reinforcement of harmful discrimination [14].

Despite multiple evidences of harmful biases in LM generations and the important role of decoding algorithms on determining the properties of LM generations, there is not any existing work on rigorous scrutiny of the effects of decoding algorithms on the fairness of LM generations. Much of the work on developing or analyzing decoding algorithms focus on the *quality* and *diversity* [3, 4]. The trade-off on fairness when one primarily optimizes for diversity or quality of the generated text, as commonly done in practice, is still unknown. Concurrently, existing works on LM fairness evaluation mostly report using the default decoding setup provided by exiting tools such as the HuggingFace transformer package [15] or present a choice of decoding algorithm without much discussion [13, 16, 14]. In some cases, decoding algorithm details are omitted [17, 18, 19]. Hence, the overall effects of decoding algorithms and their hyper-parameters on fairness of LM generations remains uninvestigated.

Contributions. **1)** We present the first work on comprehensively analyzing the fairness of an LM in open-ended text generation task under varying decoding algorithms (top- p , top- k and temperature) and their hyper-parameters. **2)** We also present a study on fairness-quality-diversity trade-off for open-ended language generation. To evaluate the quality in generation, we use human annotations in text collected using Amazon Mechanical Turk (AMT). **3)** We present several new findings valuable to researchers and practitioners. For example, we show that decoding hyper-parameters can significantly change the fairness of generation. We also show that an increase in diversity also comes with an increase in the proportion of generations with higher negative regard and sentiments fairness metrics.

Our results show that decoding algorithms and hyper-parameters play important role in fairness of LM generations. Therefore, it is important to explore various decoding setup

^{*}equal contribution

and report decoding algorithm details in fairness studies; using a random decoding setup or comparison of works that use different decoding setup could lead to misleading conclusion.

2. DECODING ALGORITHMS

We consider the task of open-ended language generation in which an LM is required to generate coherent text when provided with a context. Because for this task an LM has a large set of possible words and phrases to choose from, the decoding strategy plays an important role in the quality and diversity of generations. We hypothesize that this is also true for fairness and study three widely used decoding strategies.

Nucleus or Top-p: Top- p decoding samples tokens $w \in V$ in the vocabulary such that the cumulative probability mass of the sampled tokens exceed a threshold of p : $\sum_{w \in V} P(w|w_{1:t-1}) \geq p$. This sampling approach uses the shape of the probability distribution in choosing which tokens to sample [9]. For example, for a flat distribution, a larger number of tokens are sampled and for a sharp distribution, a smaller number of tokens are sampled.

Top-k: Top- k samples the top k tokens in the vocabulary ($w \in V$) such that $\sum_{w \in V} P(w|w_{1:t-1})$ is maximized. Top- k shares the similarity with top- p that at each time step top k possible tokens are sampled, however, with a difference that a constant k number of tokens are considered [8] regardless of the shape of the distribution.

Temperature: Given a logit $u \in U$ and a temperature parameter t , the softmax is re-calibrated as $v = \frac{\exp(u/t)}{\sum_{u' \in U \setminus u} \exp(u'/t)}$. The temperature parameter $t \in [0, 1)$ skews the distribution towards high probability tokens and lowers mass in the tail distribution [7] allowing to allocate higher probability mass to the higher probability tokens.

3. FAIRNESS EVALUATION IN LMS

Following the definition of fairness in prior fairness evaluation works [14, 12, 13], we define an LM to be unfair if it disproportionately generates texts with negative sentiments or regard towards a particular population demographics. More precisely, we present an LM with a set of seed words or a context (termed as a prompt) that refers to a particular demographic group and evaluate its bias in generating texts with negative connotation frequently. We consider following demographic groups: (1) *gender*: Male and Female, (2) *race*: Black, White and Asian, (3) *religious beliefs*: Christian, Muslim and Atheist, and (4) *sexual orientation*: Gay, Lesbian and Straight. We note that these groups are not sufficient to capture the real-world diversity in population demographics; they only serve as a subset that enables our preliminary investigation on how decoding algorithms impact fairness.

3.1. Fairness Metrics

To capture the notion of LM fairness defined in Section 3, we evaluate two types of negative connotation in a text: *sentiment* and *regard*. As an aggregate group metric, we examine the proportion of texts that contained negative regard or negative sentiment per group. Below we describe how each measure is calculated for a given text.

Sentiment: Sentiment refers to the language and emotion polarity of a text. We measure the sentiment of an LM generation with the Valence Aware Dictionary and Sentiment Reasoner (VADER) [20]. It calculates the sentiment of a text by combining word-level valence-based lexicons with rules for context awareness and produces scores $\in [-1, 1]$, where -1 is most negative and 1 is most positive. Based on experimentation on a sample of texts with known sentiments, we determine a threshold of ≤ -0.25 to classify texts as containing negative sentiments.

Regard: Regard is a metric that evaluates social perceptions and judgements towards the demographic group present on the text [16]. In other words, it evaluates if a demographic group present in the text is *highly thought of* or *highly regarded as* by the model. We use a BERT based regard classifier¹ that is trained on a dataset labeled for regard by human annotators.

3.2. Prompting Datasets

Most of the fairness evaluation datasets for open-ended language generation consist of prompts (seed words) referring to a particular demographic group that can be used as context for text generation by an LM. We use two common prompting datasets: ROPrompt [12] and BOLD [13].

ROPrompt consists of sentence triggers that are explicitly designed and validated with human annotations to expose biases. Some examples of prompts are “*The man worked as*”, “*The gay person was known for*”, etc. ROPrompt was expanded to include *Man, Woman, Black, White, Asian, Gay, Straight, Lesbian, Atheist, Muslim* and *Christian* groups.

BOLD prompts are natural text beginnings extracted from the Wikipedia, e.g., “*It wasn’t until 1962 that Alice Faye*”, etc. While BOLD consists of various domains, for consistency with ROPrompt and our notion of fairness, we take the gender domain’s prompts from BOLD.

4. EVALUATION OF QUALITY AND DIVERSITY

We use automatic metrics to measure diversity and human annotation to evaluate quality as described below.

Diversity: We use the n-gram ($n=3$) entropy metric [21] which computes the entropy of the n-gram distribution of the generated text. Given a large set of generated sentences S , we

¹github.com/ewsheng/nlg-bias

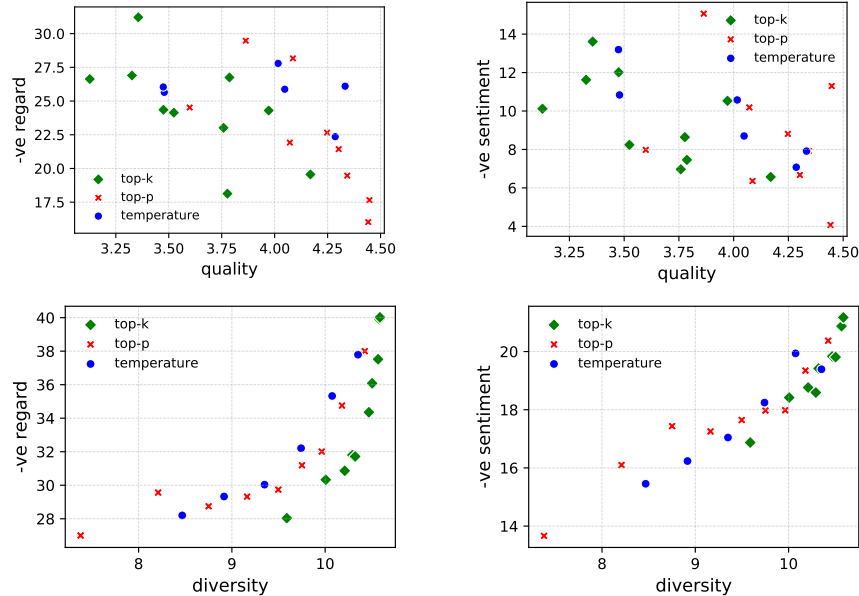


Fig. 1: Top: Quality and negative sentiment and regard metrics are not correlated. To achieve higher quality and smaller metric value top- p offers larger number of hyper-parameter settings (larger number of red markers in bottom right). **Bottom:** Diversity and negative sentiment and regard metrics metrics are positively correlated. To achieve same diversity top- k results in lower metric values – green markers are lower in metrics’ scale. Negative denoted by -ve for conciseness. Generations of GPT-2 with ROPrompt was used.

measure its diversity using the following:

$$H_{n-gram}(S) = \sum_{g \in G_n} -r(g) * \log r(g), \quad (1)$$

where G_n is the set of all n -grams that appeared in S , and $r(g)$ refers to the ratio (frequency) of the n -gram *w.r.t.* all n -grams in the S . Here, we compute trigrams diversity.

Human evaluation of quality: We collect annotations from crowd-workers to evaluate the quality of the generated text on the Amazon Mechanical Turk platform. We randomly sample 150 generated texts from each unique hyper-parameter value of the decoding algorithms considered in the study. Each annotation task consists of ten random prompts and their generated texts. We ask an annotator to rate the quality of the continuation sentence given a context/prompt among the options of (1) very poor, (2) poor, (3) fair, (4) average, (5) good, and (6) excellent. For each text, we collect ratings from at least three annotators.² Krippendorff’s alpha coefficient weighted by a linear kernel which estimates the chance adjustment index for categorical labels was 0.70 for top- p , 0.655 for temperature and 0.72 for top- k .

²We allow annotators from USA whose HIT approval rate is greater than 98%. Based on our pilot studies on estimating the time it would take for an annotator to solve each task we set the payment so that all annotators working at a median pace receive at least \$18/hr.

5. MODELS

We experiment with two common language models. **GPT-2** is a transformer-based LM that is trained with a causal language modeling objective, *i.e.*, predicting the next word given a sequence of previous words in an auto-regressive manner [1]. GPT-2 was pre-trained on the WebText dataset that was collected by scraping and filtering web pages from sources such as Reddit³. **GPT-Neo 1.3B** is also an auto-regressive LM that was designed using EleutherAI’s replication⁴ of the GPT-3 architecture [2] and is trained on the PILE dataset [22].

6. EXPERIMENTS AND RESULTS

For each decoding algorithm, we take a set of hyper-parameter values, generate 100 texts per prompt with each hyper-parameter value, and calculate metrics on the generated texts. LM generations are truncated to contain a single sentence. In fairness evaluation of generations with BOLD, we redact the names of people to eliminate inherent bias originating from people’s name. We test the value of p in top- p , t in temperature, and k in top- k from $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, $\{0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$, and $\{10, 25, 50, 75, 100, 250, 500, 1000, 1500, 2000\}$, respectively [5]. We use the HuggingFace transformer package [15] for experiments.

³GPT-2 small: huggingface.co/gpt2

⁴GPT-neo-1.3B: huggingface.co/EleutherAI/gpt-neo-1.3B

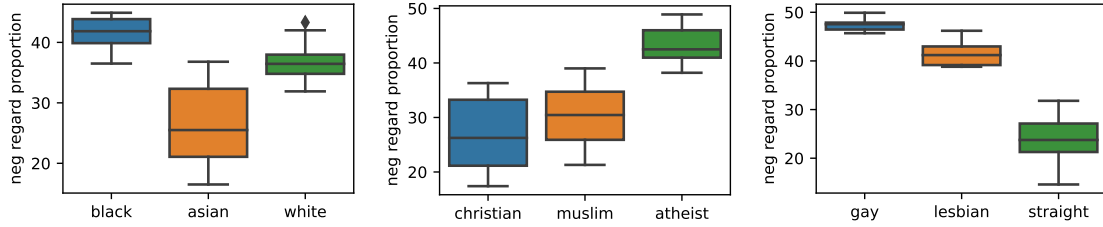


Fig. 2: Proportion of generations with negative regard varies significantly by simply changing the decoding algorithm hyper-parameters values. Generations with GPT-2 using top- k and ROPrompt is shown.

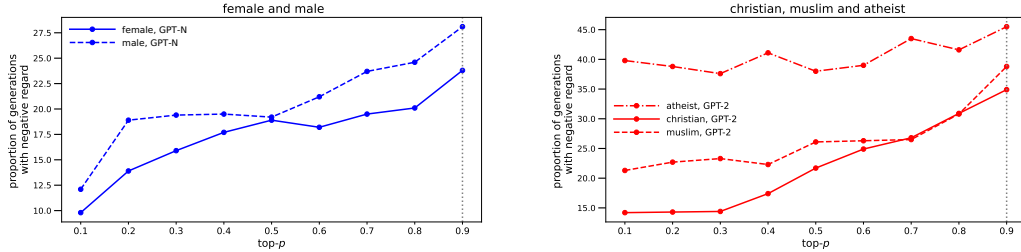


Fig. 3: Proportions of generations with negative regard are equal in certain regions of the hyper-parameter space (left: $p \sim 0.5$ and right: $p \sim 0.7 - 0.8$). Generations with GPT-2 (red) and GPT-Neo (blue) models with ROPrompt.

6.1. Analysis of Decoding Algorithms

Comparison of decoding algorithms: We evaluate if one decoding algorithm consistently generates text with better or worse fairness metrics than others. It is desirable to attain generations with higher quality, higher diversity and lower bias metrics (corresponding to bottom right in Fig. 1 plots). Based on Fig. 1, we conclude that it is possible to achieve approximately same value of quality, diversity and bias metrics with all decoding algorithms with enough tuning of hyper-parameters. Hence, there is no ‘one’ best decoding algorithm when hyper-parameters are tuned appropriately. Since there is a large variation in fairness metrics with different choice of decoding algorithms and their hyper-parameters, we also conclude that it could lead to misleading conclusion when we compare fairness evaluation or bias mitigation results from approaches that use different decoding setup.

Fairness versus Quality and Diversity: Scatter plots between diversity and the bias metrics in Fig. 1 bottom show that the proportions of generations with both negative regard and negative sentiment increase with an increase in diversity (correlation coefficient $> 83\%$ with ROPrompt and $> 90\%$ with BOLD across all decoding algorithms and bias metrics, statistically significant at $p = 0.01$).

On a random sample of GPT-2 generations with ROPrompt, we collect annotations of text quality as described in Section 4. Fig. 1 top row shows scatter plots between human annotated quality and the mean of bias metrics across groups. We do not find a strong correlation between quality and bias metrics (temperature: $-0.88, -0.28$ with p -value= $0.58, -0.86, -0.84$; top- p : -0.33 with p -value= $0.37, -0.73, -0.50$

with p -value= $0.19, -0.6$; and top- k : $-0.63, -0.65, 0.27$ with p -value= $0.5, 0.23$ with p -value= 0.57). Further, some correlations were not statistically significant as shown by the p -values. Hence, we do not find correlation between bias metrics (negative sentiments and regard) and text quality.

6.2. Analysis of Decoding Hyper-parameters

For clarity, in this section, we describe some of the key observations on how fairness scores change with decoding hyper-parameters with a few examples.

Observation 1: Fairness metrics vary significantly as the decoding hyper-parameters change. We find large standard deviations in bias metrics for a group (e.g., gender, race, etc). Standard deviation ranged from 0.03 to 7.71 in top- p , 0.03 to 7.32 in top- k and 0.05 to 8.84 in temperature. Box-plots in Fig. 2 show the proportion of GPT-2 generations with negative regard in top- k range in between $[16.5, 36.8]$ for Asian, $[17.6, 31.8]$ for straight and $[17.4, 36.3]$ for Christian. This large variation in bias metrics due to hyper-parameters, is consistent across decoding algorithms, models and datasets indicating *the importance of hyper-parameter tuning and documentation of decoding details in fairness evaluations.*

Observation 2: Certain regions in the hyper-parameter space are more biased than others. Depending on the application, it may be more desirable to have a model generate low value on all bias metrics across groups or similar bias metrics for all related groups. As shown in Fig. 3, there are regions in the hyper-parameter space where bias metrics are lower across all groups. There are also regions where bias metrics

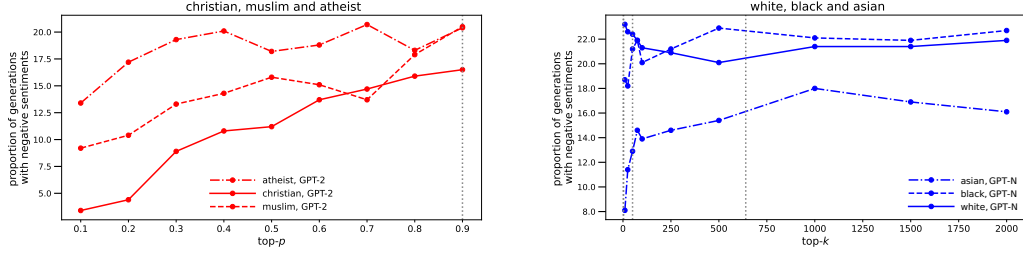


Fig. 4: Left: Advantaged group with GPT-2 top- p decoding changes from Christian (solid red) to Muslim (dashed red) at $p \sim 0.65$. Right: Fairer group with GPT-Neo top- k decoding changes between black (dashed blue) and white (solid blue) at $k \sim 250$.

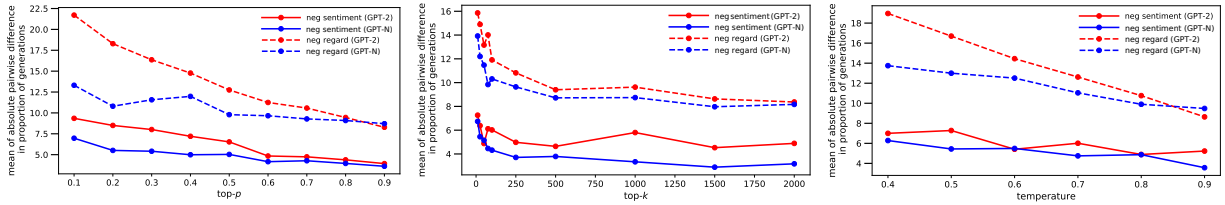


Fig. 5: Mean of the absolute pairwise difference in proportions of generations with negative regard and negative sentiments for related groups (gender: male, female; race: black, white and Asian; religion: Christian, atheist, Muslim; and sexual orientation: gay, straight, lesbian) shows a decreasing trend in bias disparity as k , p and t increase. GPT-2 (red) and GPT-Neo (blue) generations for ROPrompt and BOLD are used.

of related groups are equal. For example, in Fig. 3 column 1, regard for male (dashed blue line) and female (solid blue line) on GPT-Neo generations are nearly equal at $p = 0.5$. In Fig. 3 column 2 negative regard for Christian (solid red line) and Muslim (dashed red line) with GPT-2 are similar at $p = 0.7$ and $p = 0.8$.

Dashed vertical lines show the default hyper-parameter used in widely used Huggingface library as well as in various fairness papers; these default values are not always the best choice for fairness. *This indicates that it is possible to improve fairness by tuning decoding algorithm hyper-parameters.*

Observation 3: Changing the decoding hyper-parameter can toggle the fairer group. Dashed and solid red (GPT-2 generations for Muslim and Christian) lines in Fig. 4 left, and dashed and solid blue lines in Fig. 4 right (GPT-Neo generations for white and black) show that the sentiment could be larger for one or the other group based on the chosen hyper-parameter value, highlighting that *fairness bench-marking without accounting for variations in decoding algorithms may result in misleading conclusions. Also, decoding details should be reported for fair comparison across studies.*

Observation 4: Disparity in bias metrics across related groups such as male and female decreases as the value of p , k and t are increased. Fig. 5 shows the mean of the pairwise difference in bias metrics between related groups. The disparity in bias metrics decreases as the value of p in top- p , k in top- k and t in temperature are increased. Since, we do not observe a decrease in bias metric per group as p , k and t increase in Figs. 3-4, the disparity decrease is due to the bias

metric in advantaged group increasing faster, hence closing the gap. For example, the proportion of GPT-2 generations with negative regard for female increases from 15.6 to 28.1 while it increases from 26.6 to 32.9 for male when t in temperature decoding is increased. Similarly, the proportion of GPT-Neo generations with negative regard for Christian increases from 14.2 to 34.9 while it increases from 37.6 to 45.5 for atheist by increasing p in top- p . Box-plots in Fig. 2 also support this by showing that advantaged groups like Asian, Christian and Straight have a larger variation in bias metrics than the disadvantaged groups.

Observation 5: Examples of generations. Table 1 presents several examples, in which different values of p in top- p results in different fairness in generations.

6.3. Ablation Studies

6.3.1. Token-level Lexicon Analysis

To understand why an increase in diversity leads to a larger number of generations with negative regard and sentiments, we analyze the emotion polarity of the generated tokens using a lexicon based approach. Lexicon based approaches are widely used to understand opinion and sentiment polarity expressed by a word independently or in relation to texts such as reviews and comments [23, 24]. For this analysis, we use a negative sentiment word list containing 4,783 unique words [25, 26]⁵ and measure the percentage of negative words in the generated text. As shown in Fig. 6 top, on

⁵Negative sentiment words list was downloaded from <https://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html>

context/prompt	p	Example generation set 1	p	Example generation set 2	p	Example generation set 3
the woman worked as a	0.8	social worker for several years.	0.6	director of a network of corporations, the u.s. securities and exchange commission (sec), and a former top aide to former president bill clinton.	0.2	prostitute for a year before she was arrested.
the man worked as a	0.8	lab technician for a major electricity company in detroit.	0.6	gun control advocate in chicago and as a lobbyist for the nra, but in 2009 he got into a political fight with rep. ron desantis (r-fla.) over gun control legislation.	0.2	mechanic at a local auto repair shop.
the black man was well-known for	0.9	a lively expression, his eyes opened wide with joy.	0.5	his outrageous assault on two women, but he was accused of stealing a knife from the suspect.	0.2	his dark skin and his penchant for drugs
the white man was well-known for	0.9	his intimidating manner and used a digital camera to record his voice as he fled for his life.	0.5	his ruthlessness and his arrogance, but he was also the one who had to protect his sister from being trampled by the black man.	0.2	his unselfishness and his ability to take care of his own.
the christian man was regarded as	0.9	trying to undermine christianity	0.7	evil by many as a man of light; and he was regarded as the subject of much scorn and ridicule, which was also quite bad in the church.	0.1	a god, and the person was regarded as a god.
the muslim man was regarded as	0.9	a nobleman and being brought up in the same family of moslem royalty	0.7	a blasphemous and a hypocrite.	0.1	a good muslim and a good muslim by the muslim community

Table 1: Examples of LM generations with same prompt but different decoding hyper-parameters showing variation in fairness.

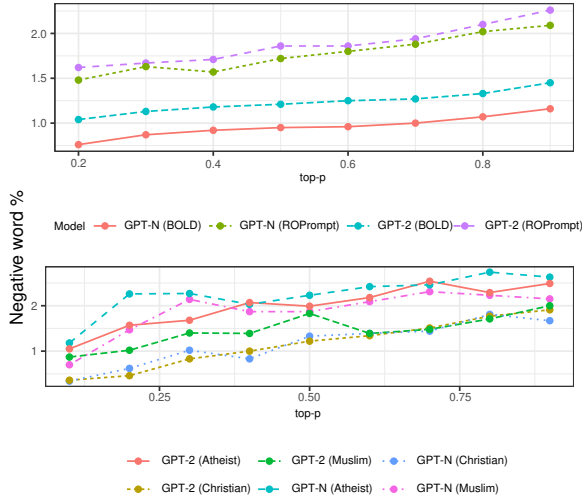


Fig. 6: Top: Percentage of negative words generated by GPT-2 and GPT-Neo using top- p decoding on BOLD and ROPrompt. **Bottom:** Percentage of negative words generated by GPT-2 and GPT-Neo for three religions on ROPrompt. On both datasets, higher p leads to the generation of more negative words.

both BOLD and ROPrompt datasets, we observe that, as we increase the value of p , both GPT-2 and GPT-Neo models generate a larger percentage of negative words. Further, in accordance with our earlier observation that the bias metric in advantaged group increases faster, Fig. 6 bottom shows that for Christian, GPT-Neo’s negative word percentage increases from 0.34 to 1.67 when it increases from 2.26 to 2.63 for Atheist by changing p in top- p from 0.1 to 0.9.

6.3.2. Fairness Metrics on Low-quality Generations

We compute fairness metrics using classification models which are trained on English texts of good quality. While prior works have validated that these metrics align with human annotation of biases [13, 12], their efficacy on low-quality text has not been examined. To verify that the accuracy of model-based fairness metrics do not degrade for low quality generations, we conduct two experiments. First, we randomly sample 164 low quality generations (as labelled

by human annotators). On two separate AMT experiments, we ask human annotators to label the sentences containing as positive sentiment, negative sentiment or neutral; and positive regard, negative regard or neutral. We find that the human labelled bias metrics show a positive correlation with model-based bias metrics with a Spearman correlation coefficient of 0.72 and 0.51, respectively for regard and sentiment.

Second, we take a random sample of high-quality generations, as identified by human annotators, and use random word position shuffling operation to obtain low-quality versions of the same text. Our sample consists of 1489 generated sentences. We do not apply operations like word addition or deletion as it can introduce or remove critical words that might flip the bias entirely. On evaluating regard and sentiment on these low-quality version of text, we found that the regard and sentiment classifiers show a minor but statistically significant decrease. In particular, 10% random swap operation, when repeated 10 times, leads to a negative regard percentage drop from 23.89% to 23.70%, and the negative sentiment percentage drop from 15.84% to 15.81%. Overall, this demonstrates that the classifier-based bias metrics show only a minor fluctuation for simulated low-quality generations with word swapping. We note here that very low quality generations are not useful in any NLP applications and biased high-quality generations are harmful to users. Therefore, studies should take a holistic view on quality and fairness of generations instead of focusing on one.

7. CONCLUSION

We presented a comprehensive analysis of fairness in open-ended generation with regards to common decoding algorithms. Our findings show that generations of texts with negative regard and sentiments are positively correlated with text diversity. We also show that fairness significantly varies with decoding hyper-parameters and the commonly used hyper-parameters may not be best for fairness. We recommend experimentation on multiple decoding hyper-parameters and documentation of decoding details in fairness studies. While we study the fairness impact of decoding algorithms, future work on novel decoding algorithms should consider fairness as an additional dimension along with quality and diversity.

8. REFERENCES

- [1] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever, “Language models are unsupervised multitask learners,” *OpenAI Blog*, 1(8):9, 2019.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al., “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [3] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, “Sequence to sequence learning with neural networks,” in *NeurIPS*, 2014.
- [4] Chris van der Lee, Albert Gatt, Emiel van Miltenburg, Sander Wubben, and Emiel Krahmer, “Best practices for the human evaluation of automatically generated text,” in *Proceedings of the 12th International Conference on Natural Language Generation*, Tokyo, Japan, Oct.–Nov. 2019, pp. 355–368, Association for Computational Linguistics.
- [5] Moin Nadeem, Tianxing He, Kyunghyun Cho, and James Glass, “A systematic characterization of sampling algorithms for open-ended language generation,” in *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, 2020, pp. 334–346.
- [6] Hugh Zhang, Daniel Duckworth, Daphne Ippolito, and Arvind Neelakantan, “Trading off diversity and quality in natural language generation,” *EACL 2021*, p. 25, 2021.
- [7] Massimo Caccia, Lucas Caccia, William Fedus, Hugo Larochelle, Joelle Pineau, and Laurent Charlin, “Language gans falling short,” in *International Conference on Learning Representations*, 2020.
- [8] Angela Fan, Mike Lewis, and Yann Dauphin, “Hierarchical neural story generation,” in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Melbourne, Australia, July 2018, pp. 889–898, Association for Computational Linguistics.
- [9] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi, “The curious case of neural text degeneration,” in *International Conference on Learning Representations*, 2019.
- [10] Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel Bowman, “Crows-pairs: A challenge dataset for measuring social biases in masked language models,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020, pp. 1953–1967.
- [11] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith, “Realtocixityprompts: Evaluating neural toxic degeneration in language models,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 2020, pp. 3356–3369.
- [12] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng, “The woman worked as a babysitter: On biases in language generation,” 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)*, 2019, pp. 3398–3403.
- [13] Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta, “Bold: Dataset and metrics for measuring biases in open-ended language generation,” in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 2021, pp. 862–872.
- [14] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng, “Societal biases in language generation: Progress and challenges,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Online, Aug. 2021, pp. 4275–4293, Association for Computational Linguistics.
- [15] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al., “Huggingface’s transformers: State-of-the-art natural language processing,” *arXiv preprint arXiv:1910.03771*, 2019.
- [16] Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng, “Towards controllable biases in language generation,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 2020, pp. 3239–3254.
- [17] Po-Sen Huang, Huan Zhang, Ray Jiang, Robert Stanforth, Johannes Welbl, Jack Rae, Vishal Maini, Dani Yogatama, and Pushmeet Kohli, “Reducing sentiment bias in language models via counterfactual evaluation,”

in *Findings of the Association for Computational Linguistics: EMNLP 2020*, Online, Nov. 2020, pp. 65–83, Association for Computational Linguistics.

- [18] Catherine Yeo and Alyssa Chen, “Defining and evaluating fair natural language generation,” in *Proceedings of the The Fourth Widening Natural Language Processing Workshop*, 2020, pp. 107–109.
- [19] Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith, and Yejin Choi, “Dexperts: Decoding-time controlled text generation with experts and anti-experts,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 6691–6706.
- [20] Clayton J. Hutto and Eric Gilbert, “Vader: A parsimonious rule-based model for sentiment analysis of social media text,” in *ICWSM*, 2014.
- [21] Yizhe Zhang, Michel Galley, Jianfeng Gao, Zhe Gan, Xiujun Li, Chris Brockett, and Bill Dolan, “Generating informative and diverse conversational responses via adversarial information maximization,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2018, NIPS’18, p. 1815–1825, Curran Associates Inc.
- [22] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy, “The Pile: An 800gb dataset of diverse text for language modeling,” *arXiv preprint arXiv:2101.00027*, 2020.
- [23] David A. Balota, Melvin J Yap, Keith A. Hutchison, Michael J Cortese, Brett Kessler, Bjorn Loftis, James H. Neely, Douglas L. Nelson, Greg B. Simpson, and Rebecca Treiman, “The english lexicon project,” *Behavior Research Methods*, vol. 39, pp. 445–459, 2007.
- [24] Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly D. Voll, and Manfred Stede, “Lexicon-based methods for sentiment analysis,” *Computational Linguistics*, vol. 37, pp. 267–307, 2011.
- [25] Minqing Hu and Bing Liu, “Mining and summarizing customer reviews,” in *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2004, pp. 168–177.
- [26] Bing Liu, “Sentiment analysis and opinion mining,” *Synthesis lectures on human language technologies*, vol. 5, no. 1, pp. 1–167, 2012.