

Detect and Perturb: Neutral Rewriting of Biased and Sensitive Text via Gradient-based Decoding

Zexue He, Bodhisattwa Prasad Majumder, Julian McAuley
Department of Computer Science and Engineering, UC San Diego
{zehe, bmajumde, jmcauley}@eng.ucsd.edu

Abstract

Written language carries explicit and implicit biases that can distract from meaningful signals. For example, letters of reference may describe male and female candidates differently, or their writing style may indirectly reveal demographic characteristics. At best, such biases distract from the meaningful content of the text; at worst they can lead to unfair outcomes. We investigate the challenge of re-generating input sentences to ‘neutralize’ sensitive attributes while maintaining the semantic meaning of the original text (e.g. is the candidate qualified?). We propose a gradient-based rewriting framework, **Detect and Perturb to Neutralize (DEPEN)**, that first detects sensitive components and masks them for regeneration, then perturbs the generation model at decoding time under a *neutralizing* constraint that pushes the (predicted) distribution of sensitive attributes towards a uniform distribution. Our experiments in two different scenarios show that DEPEN can regenerate fluent alternatives that are neutral in the sensitive attribute while maintaining the semantics of other attributes.

1 Introduction

Language data often carries implicit biases or contains sensitive information that may have negative consequences for human and machine understanding. For example, a person’s choice of vocabulary can reveal their social identity (age, gender, or political affiliation) (Nguyen et al., 2013); a few examples are shown in Table 1. Such information can potentially bias machine predictions as well as human judgment, leading to unfair outcomes.

Hiding sensitive information in textual data—including text that carries *implicit* bias—is an essential task. In this paper we consider the setting of graduate school admissions as a case-study, where fair evaluation of applicants should depend on academic performance or research potential, irrespective of nationality, gender, etc. Text from reference

Text	Attr.
1. She gone dance without da bands lol.	Race
2. Hahaahahaha wwatching rtl gemist hol-land, bigga is cryingg it’s killingg me.	Age
3. Tasted as amazing as the first sip I took! Definitely would recommend	Gender
4. PERSON-B-1 is adorable with pleasant and easy-going personality.	Gender

Table 1: Examples of scenarios that reveal sensitive attributes (Attr.). Highlighted words are markers of such sensitive information. Example 1 shows an excerpt of a tweet written by an African-American revealed by vocabulary usage (future tense of *gone* → “is going to”) (Blodgett et al., 2018). Example 2 is a tweet from a young person (Nguyen et al., 2013). Example 3 is a review by a female (from Yelp dataset (Reddy and Knight, 2016)) while Example 4 describes a female applicant in a graduate admission reference letter (our data).

letters is colored by many biases: letter writers may (possibly unintentionally) write about male and female candidates differently, or may use language that reflects their (the writer’s or the applicant’s) cultural background. Eliminating these attributes from the decision making process is challenging because (1) the sensitive information is often implicit and confounded with other attributes, and (2) a parallel corpus with *unbiased* text is not available.

Based on these motivations, we define our task as: given an input sentence associated with both meaningful and sensitive attributes (e.g. a discussion of a female student’s research potential), *re-generate* the input in a way that *neutralizes* one or many *sensitive attributes* with *minimal edits*, i.e., so as to maintain the fluency, coherency, and semantic meaning of the original sentence.

To this end, we propose a gradient-based decoding framework for text re-generation by neutralizing a sensitive attribute: **Detect and Perturb to Neutralize (DEPEN)**. We realize the framework in two steps (Figure 1). First we automatically detect the parts of the input sentence that reveal the sensitive attribute, and mask them; while this can be as

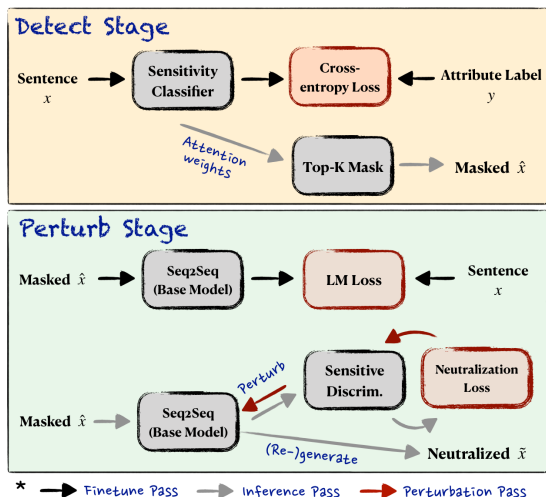


Figure 1: The dataflow of DEPEN. Details of the Detect stage and Perturb stage are explained in Section 2.

simple as a gendered pronoun (‘he/she’), we find many cases where choices of adjectives or phrasing are associated with group identity. Second, we regenerate a complete sentence from the unmasked part of the input so that the output no longer reveals the sensitive attribute. We do this by perturbing the final hidden states of a conditional language model that is finetuned to generate a complete sentence from masked tokens. Perturbation is done to modify the hidden states in a ‘neutral’ (i.e., so that the hidden state cannot predict the sensitive attribute) direction while maintaining fluency and semantic meaning. We conduct two experiments to show that DEPEN generalizes across scenarios. We first experiment with a Graduate Admissions Reference letter dataset where DEPEN rewrites the sentences from a letter to neutralize attributes such as gender or nationality. So that we can release a reproducible benchmark, we also experiment with *Goodreads* review data (Wan and McAuley, 2018); here we treat genres as a sensitive attribute (i.e., maintain the essence of a review without revealing the genre).

2 DEPEN

As shown in Figure 1, our neutralizing approach DEPEN¹ has two stages: Detect and Perturb.

2.1 Detect: mask the sensitive parts

First we detect parts of the original input sentence x that are predictors of the target sensitive attribute \mathcal{A} . Suppose we have a corpus containing N documents and their associated label y for \mathcal{A} ; we train a classifier f_θ to

minimize $\frac{1}{M} \sum_{i=1}^N \sum_{j=1}^{|X^i|} \mathcal{L}(f(x_j^i; \theta), y^i)$, where X^i is the i -th document and x_j^i is the j -th sentence, M is the number of sentences, and \mathcal{L} is the cross-entropy loss for classifying sensitive attributes.

Following Jain et al. (2020), we take self-attention scores of all input tokens w.r.t. the [CLS] token (Devlin et al., 2018) from the final hidden layers and normalize them to measure how salient each token is for predicting \mathcal{A} . We use BERT as the attribute classifier f .

Next, we mask the top- $k\%$ (k is a hyperparameter) salient tokens to obtain the intermediate output as \hat{x}_j^i that does not contain any significant predictor of \mathcal{A} according to f .

2.2 Perturb to Neutralize

To regenerate a neutral version \tilde{x} of the original input sentence x we need a generative model that can reconstruct a sentence from the unmasked tokens. For this we train a sequence-to-sequence (Seq2Seq) model that takes \hat{x}_j^i as input and x_j^i as output. We finetune a BART model as our base Seq2Seq model g . Ideally, we want g to regenerate a version that remains neutral to the attribute \mathcal{A} . But since we do not have attribute-neutral ground-truth, we cannot guarantee that inference from g will hold attribute neutrality. Hence, we guide g using a gradient-based inference method so that the regenerated output remains attribute-neutral. We are inspired by PPLM (Dathathri et al., 2019) that introduced gradient-based inference from transformer-based language models. Similar inference-time perturbation approaches also have been proposed for applications such as clarification question generation (Majumder et al., 2021b) and dialog generation (Majumder et al., 2021a).

PPLM primarily performs gradient-based decoding that encourages the generation to maintain fluency according to the base autoregressive generative model while honoring a discriminative constraint, such as maintaining a particular attribute. In our work, we modify PPLM to accommodate a new decoding constraint for achieving neutrality. We also adapt a Seq2Seq transformer model as a base model to perform autoregressive inference using PPLM-style gradient decoding.

Generate with Neutralizing Constraints Contrary to PPLM, which boosts the log-likelihood (LL) of a certain attribute, our case requires the generation is *neutral* toward an attribute (e.g. the text should be neither ‘female’ nor ‘male’). Since

¹<https://github.com/ZexueHe/DEPEN>.

we do not have explicit labels for neutrality, we modify our decoding constraint in the following.

Suppose there are $|\mathcal{C}|$ categories for \mathcal{A} and we want to re-generate a sentence \tilde{x}_j^i which minimizes the KL-divergence between a uniform distribution over \mathcal{C} and the discriminative distribution of the sensitive attribute \mathcal{A} . We define it as our neutralization constraint $\mathcal{L}_{\text{ntrl}}$

$$\begin{aligned} & \arg \min_{\tilde{x}_j^i} D_{KL} (U(\mathcal{C}) \parallel p(y^i | \tilde{x}_j^i)) \\ &= \arg \min_{\tilde{x}_j^i} H(U(\mathcal{C}), p(y^i | \tilde{x}_j^i)) - H(U(\mathcal{C})) \\ &= \arg \min_{\tilde{x}_j^i} \underbrace{- \sum_{a \in \mathcal{C}} \frac{1}{|\mathcal{C}|} \log p(y^i = a | \tilde{x}_j^i)}_{\mathcal{L}_{\text{ntrl}}} \end{aligned}$$

where $H(\cdot)$ is the entropy and $U(\cdot)$ denotes the uniform distribution.

Since ground truth is not available, we resort to an unsupervised decoding technique using the left-to-right decoder from the Seq2Seq model. During inference, we keep the encoder of the base model g fixed while perturbing the hidden states of the decoder. A gradient w.r.t. the neutralization loss $\mathcal{L}_{\text{ntrl}}$ shifts the hidden state representations toward neutrality during backpropagation. To realize the effect of backward gradient updates, we accumulate gradients for multiple passes and then update the hidden representations. Once we update decoder hidden states, a forward pass is made to maintain the fluency of the base language model. Backward and forward pass alternate until we see the desired neutralization effect in the generated text.

3 Experiments

3.1 Datasets

Reference letters a real-world dataset of students considered for admission to a graduate program of a large US university,² containing applicant profiles including reference letters, binary gender information, nationality, and a binary admission decisions. We consider 18,865 applicants with 29,170 reference letters, among which 22,201 letters are used for training classifiers and 6,969 for testing or rewriting. We conduct two experiments with gender and nationality (processed to be 4 dominant classes) as sensitive attributes separately, and use admission decisions as the outcome for further evaluating whether the ‘signal’ is preserved.

²Our investigation is IRB-approved. Details are anonymized even in our private version.

GoodReads a book review dataset (Wan and McAuley, 2018) containing user reviews, star ratings, and genres. We randomly sample 3000 reviews each from the *Children’s* and *Mystery* genres. We use 5000 reviews for training and the rest for testing. We define the binary genre as the sensitive attribute, and quantize ratings to three levels (positive, negative, neutral) as the outcome.

3.2 Evaluation Metrics

Bias: We use the accuracy (Acc.) and confidence (Conf.) of a sensitive classifier to evaluate bias. **Fluency:** We use the Pseudo Log-Likelihood (PLL) of Salazar et al. (2020) to measure the fluency of our generated model. **Coherence:** We use the BLEU4 score of the generated sentence w.r.t. its input and accuracy of an outcome (Out.) classifier to measure how much content is maintained.

3.3 Baseline Models

We evaluate four debiasing approaches (all of which generate without parallel ground truth) and two variants of DEPEN as baselines:

- Rule-based (RB): replace words with rules (e.g. *he/she* → *they*, see Appendix A.1).
- Weighed Decoding (WD): a decoding method (Ghazvininejad et al., 2017) by reducing the generation probability of detected sensitive tokens to a hyperparameter α (we set $\alpha = 0.2$).
- Adversarial Training (ADV): a Seq2Seq autoencoder with a gradient reversal layer (Ganin and Lempitsky, 2015) that propagates gradients of the sensitive discriminator to the encoder.
- Privacy-Aware Text Rewriting (PATR): we reimplement the adversarial back-translation rewriting model of Xu et al. (2019).
- DEN: DEPEN w/o Perturb, generates \tilde{x} from \hat{x} with the finetuned base model g .
- PEN: DEPEN w/o Detect, generates \tilde{x} from x by neutrally perturbing a normal Seq2Seq.

3.4 Results and Analysis

Results are shown in Table 2. For debiasing metrics, DEPEN leads to a decrease (as desired) in Acc. and Conf. to around 0.5 for all experiments. We note that PEN generates sentences with a normal BART designed for common Seq2Seq tasks like summarization or translation, so in spite of a somewhat better accuracy drop, regenerated sentences

Model	CS Admission Dataset										GoodReads Dataset				
	Gender (binary)					Nationality (4 classes)					Genre (binary)				
	Acc.	Conf.	PLL	BLEU4	Out.	Acc.	Conf.	PLL	BLEU4	Out.	Acc.	Conf.	PLL	BLEU4	Out.
Original	0.9247	0.9002	-4.8134	1.0000	0.6321	0.7487	0.6660	-4.6511	1.000	0.6741	0.7557	0.7165	-4.3154	1.0000	0.6551
RB	0.7397	0.6614	-5.0973	0.8761	0.6333	0.7470	0.6665	-4.8624	0.9974	0.6353	0.7297	0.6938	-4.4106	0.9699	0.6543
WD	0.7125	0.6940	-5.0520	0.3781	0.6101	0.6303	0.5568	-4.6771	0.5251	0.6105	0.6565	0.6885	-4.5162	0.2571	0.5905
ADV	0.9197	0.8970	-5.9049	0.3979	0.5818	0.7091	0.6302	-5.8053	0.3838	0.5551	0.7364	0.7149	-4.7013	0.2917	0.5978
PATR	0.8797	0.8528	-5.0034	0.5278	0.6071	0.7148	0.629	-4.7511	0.5336	0.6154	0.7451	0.7077	-4.4254	0.3637	0.5979
DEN	0.7546	0.7375	-4.8695	0.4449	0.6171	0.6416	0.5694	-4.6696	0.5261	0.5818	0.6815	0.6534	-3.9767	0.3749	0.6015
PEN	0.5002	0.4617	-5.1048	0.0825	0.5871	0.2486	0.2489	-5.0260	0.0652	0.6245	0.5362	0.5219	-4.9699	0.2471	0.5728
DEPEN	0.5157	0.4935	-4.8464	0.6356	0.6298	0.5242	0.4564	-4.6623	0.5357	0.6491	0.5915	0.5665	-4.3798	0.3747	0.6168

Table 2: Results on Reference Letters and GoodReads data (see Section 3.4).

differ vastly from inputs, which can be seen from low BLEU4 scores (0.0825 for gender and 0.06 for nationality). WD also lowers bias, but it can abruptly interrupt the generation by reducing the probabilities of certain (sensitive) tokens affecting the overall language model fluency.

We also report the accuracy of predicting outcome variables (Out.), i.e., admission decisions or review sentiment (which are *not* used for training).

For fluency DEN has the highest (i.e., best) PLL but fails to debias (high Acc. and Conf.). DEPEN maintains high fluency while also debiasing.

RB has the highest coherence, though we find that regenerated sentences are extremely similar to the input (with many biased terms persisting) due to simple replacement rules. RB has extremely high BLEU4 scores (0.9974 for nationality and 0.9699 for GoodReads). PATR also demonstrates its effectiveness on language quality (fluency and coherence) due to the paraphrasing capability of back-translation, however it fails to debias well as it still shows high Acc. and Conf. in bias classification (more in Appendix A.2).

DEPEN beats the baselines by achieving a balance across bias mitigation, fluency, and coherency, and fidelity w.r.t. the predicted outcome. Manual inspection revealed that automatic metrics are suggestive of how humans perceive neutrality.

3.5 Case Study

We provide an example in Table 3, in which a referrer comments on the mock classes of a student. More examples and findings are shown in Appendix A.2. Besides the obvious gendered indicators *Her/girl*, the words *lovely* and *popular* are also considered as gender-predictive. For RB, such adjectives strain the ability of humans to design perfect rules, not only because it is hard to enumerate all such words but also due to their context-dependence (e.g. ‘elegant’ may carry different bias

Model	Re-generated
Original	Her course really attracted others, that made this lovely girl really popular in classroom.
RB	<i>Their</i> course really attracted others, that made this <i>lovely person</i> really <i>popular</i> in classroom.
WD	<i>The</i> course instantly attracted others, that made this <i>young man</i> really <i>active</i> in classroom.
ADV	<i>Her</i> course really attracted others, that made this <i>excited girl</i> really <i>popular</i> in classroom.
PATR	<i>Her</i> course really attracted others, which made this <i>lovely girl</i> really <i>popular</i> in class.
DEN	<i>The</i> course instantly attracted others, that made this <i>young man</i> really <i>active</i> in classroom
PEN	A class almost one third time I got on the topic, but it’s true for the classroom at home.
DEPEN	<i>This</i> course instantly attracted others, that made this <i>young student</i> really <i>shine</i> in classroom.

Table 3: Re-generated examples. We show detected sensitive words in red, and edited words in italics.

if it describes a student versus a student’s theorem). Simple replacement (e.g. *their*) also yields ungrammatical sentences. For WD and DEN, without a neutralization constraint, they select candidates that satisfy the language model, but may choose (e.g.) *man*, leading to no reduction in attribute sensitivity, and (e.g.) *active* which changes the semantic meaning. As a black-box rewriting method with strong reconstruction signals, it’s harder to control ADV to meet all expectations simultaneously. PATR also fails to debias. However, DEPEN can edit the sensitive parts while maintaining fluency and semantic meaning.

4 Related Work

Debiasing Language Generation There are three main streams to debias NLG tasks: counterfactual data augmentation (Lu et al., 2020; Chen et al., 2018); training-time methods (Huang et al., 2020; Liu et al., 2020a,b; Kaneko and Bollegala, 2021; Pryzant et al., 2020); and inference-time methods. Saunders and Byrne (2020) mitigate gender bias in machine translation via transfer learning using handcrafted gender-balanced datasets. Sheng et al. (2020) generate with well-formulated bias

triggers based on (Wallace et al., 2019) to equalize biases between demographics. Dathathri et al. (2019) propose a gradient-based method for controllable generation and show its efficacy in toxicity reduction. However, all these methods require explicit labels or parallel data regarding the desired attribute.

Re-writing Here specific parts of the original text are revised to be more aligned with a target attribute (Thompson, 2013). Representative approaches use an encoder-decoder setup with a discriminator (e.g. style) (Romanov et al., 2018; Dai et al., 2019; John et al., 2019; Aho and Ullman, 1972; Majumder et al., 2021a,b), backtranslation (Lample et al., 2018; Prabhume et al., 2018; Xu et al., 2019), pretraining (Duan et al., 2020; Zhou et al., 2021), or use retrieval framework (Sudhakar et al., 2019). A few approaches adapt these techniques for debiasing. Zmigrod et al. (2019) mitigate gender bias by converting between masculine- and feminine-inflected sentences with data augmentation; Ma et al. (2020) jointly train a reconstruction and an out-of-domain paraphrasing task to correct bias, which requires a parallel corpus with attribute-sensitive (e.g. gender) verbs assigned and masked. In contrast, we aim to rewrite neutrally without human guidance.

5 Conclusion

In this work, we propose a gradient-based rewriting framework, DEPEN, to neutralize a text that carries sensitive information (e.g., gender) by detecting the sensitive-predictable parts and perturbing the regeneration via a neutralization constraint. The constraint will shift the re-generated sentences to be uniform distributed for the sensitive attribute (e.g., neither male nor female) with minimal editing to maintain the semantic content.

Acknowledgments

We thank Taylor Berg-Kirkpatrick, Jianmo Ni, Yuheng Zhi, and anonymous reviewers for their valuable suggestions to this work. Our Reference Letter dataset is built from anonymized admission data of UC San Diego’s CSE department, and its use is supported by our IRB. BPM is partly supported by a Qualcomm Innovation Fellowship and NSF Award #1750063.

Ethical considerations

While a debiasing system is intended to mitigate fairness issues in natural language, such a system could certainly have unwanted side effects. Most critically, removing bias may to some extent eliminate meaningful signal from the data, or subtly alter the intended meaning of a sentence. A malicious user could adversarially maximize the neutralization constraint which would result in enhancing the bias in the input sentence. A system like ours should likely not be used as a ‘black box,’ but would best be used in a setting where its outputs can be ‘audited’ to ensure that semantic meaning is preserved, e.g. by a letter writer trying to improve their own writing or by a neutral third party.

References

- Alfred V. Aho and Jeffrey D. Ullman. 1972. *The Theory of Parsing, Translation and Compiling*, volume 1. Prentice-Hall, Englewood Cliffs, NJ.
- Su Lin Blodgett, Johnny Wei, and Brendan O’Connor. 2018. Twitter universal dependency parsing for african-american and mainstream american english. In *ACL*, pages 1415–1425.
- Wei-Fan Chen, Henning Wachsmuth, Khalid Al Khatib, and Benno Stein. 2018. Learning to flip the bias of news headlines. In *INLG*, pages 79–88.
- Ning Dai, Jianze Liang, Xipeng Qiu, and Xuan-Jing Huang. 2019. Style transformer: Unpaired text style transfer without disentangled latent representation. In *ACL*, pages 5997–6007.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019. Plug and play language models: A simple approach to controlled text generation. In *ICLR*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Yu Duan, Canwen Xu, Jiaxin Pei, Jialong Han, and Chenliang Li. 2020. Pre-train and plug-in: Flexible conditional text generation with variational auto-encoders. In *ACL*.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. Hafez: an interactive poetry generation system. In *ACL, System Demonstrations*, pages 43–48.

- Po-Sen Huang, Huan Zhang, Ray Jiang, Robert Stanforth, Johannes Welbl, Jack Rae, Vishal Maini, Dani Yogatama, and Pushmeet Kohli. 2020. Reducing sentiment bias in language models via counterfactual evaluation. In *EMNLP: Findings*, pages 65–83.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C Wallace. 2020. Learning to faithfully rationalize by construction. In *ACL*, pages 4459–4473.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled representation learning for non-parallel text style transfer. In *ACL*, pages 424–434.
- Masahiro Kaneko and Danushka Bollegala. 2021. Debiasing pre-trained contextualised embeddings. In *ACL*, pages 1256–1266.
- Guillaume Lample, Sandeep Subramanian, Eric Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2018. Multiple-attribute text rewriting. In *ICLR*.
- Haochen Liu, Jamell Dacon, Wenqi Fan, Hui Liu, Zitao Liu, and Jiliang Tang. 2020a. Does gender matter? towards fairness in dialogue systems. In *COLING*, pages 4403–4416.
- Haochen Liu, Wentao Wang, Yiqi Wang, Hui Liu, Zitao Liu, and Jiliang Tang. 2020b. Mitigating gender bias for neural dialogue generation with adversarial learning. In *EMNLP*, pages 893–903.
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2020. Gender bias in neural natural language processing. In *Logic, Language, and Security*, pages 189–202. Springer.
- Xinyao Ma, Maarten Sap, Hannah Rashkin, and Yejin Choi. 2020. Powertransformer: Unsupervised controllable revision for biased language correction. In *EMNLP*, pages 7426–7441.
- Bodhisattwa Prasad Majumder, Taylor Berg-Kirkpatrick, Julian J. McAuley, and Harsh Jhamtani. 2021a. Unsupervised enrichment of persona-grounded dialog with background stories. In *ACL*, pages 585–592.
- Bodhisattwa Prasad Majumder, Sudha Rao, Michel Galley, and Julian J. McAuley. 2021b. Ask what’s missing and what’s useful: Improving clarification question generation using global knowledge. In *NAACL-HLT*, pages 4300–4312.
- Dong Nguyen, Rilana Gravel, Dolf Trieschnigg, and Theo Meder. 2013. How old do you think i am? a study of language and age in twitter. In *ICWSM*, volume 7.
- Shrimai Prabhunoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. 2018. Style transfer through back-translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 866–876.
- Reid Pryzant, Richard Diehl Martinez, Nathan Dass, Sadao Kurohashi, Dan Jurafsky, and Diyi Yang. 2020. Automatically neutralizing subjective bias in text. In *AAAI*, volume 34, pages 480–489.
- Sravana Reddy and Kevin Knight. 2016. Obfuscating gender in social media writing. In *Workshop on NLP and Computational Social Science*, pages 17–26.
- Alexey Romanov, Anna Rumshisky, Anna Rogers, and David Donahue. 2018. Adversarial decomposition of text representation. *arXiv preprint arXiv:1808.09042*.
- Julian Salazar, Davis Liang, Toan Q Nguyen, and Katrin Kirchhoff. 2020. Masked language model scoring. In *ACL*, pages 2699–2712.
- Danielle Saunders and Bill Byrne. 2020. Reducing gender bias in neural machine translation as a domain adaptation problem. In *ACL*, pages 7724–7736.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2020. Towards controllable biases in language generation. In *EMNLP: Findings*, pages 3239–3254.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. 2019. “transforming” delete, retrieve, generate approach for controlled text style transfer. In *EMNLP*, pages 3269–3279.
- Travis Thompson. 2013. Autism research and services for young children: History, progress and challenges. *Journal of Applied Research in Intellectual Disabilities*, 26(2):81–107.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for attacking and analyzing nlp. In *EMNLP-IJCNLP*, pages 2153–2162.
- Mengting Wan and Julian McAuley. 2018. Item recommendation on monotonic behavior chains. In *RecSys*, pages 86–94.
- Qionghai Xu, Lizhen Qu, Chenchen Xu, and Ran Cui. 2019. Privacy-aware text rewriting. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 247–257.
- Wangchunshu Zhou, Tao Ge, Canwen Xu, Ke Xu, and Furu Wei. 2021. Improving sequence-to-sequence pre-training via sequence span rewriting. In *EMNLP*.
- Ran Zmigrod, Sabrina J Mielke, Hanna Wallach, and Ryan Cotterell. 2019. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. In *ACL*, pages 1651–1661.

A Appendix

A.1 Details about Baselines

Rule-based Model Detailed rules are described in Table 4. For gender, we follow the handout³ for mitigation. For nationality, though we have masked the sensitive information with Named Entity Recognition (NER), there are a few cases where NER fails, such as “Chinese Mathematical Olympiad”; to handle this we delete a list of country/city/nationality names. Since we can’t precisely formulate the special patterns corresponding to applicants from different nationalities, we count unique bi-grams in the top-100 bi-gram list of each category as additional rules. For GoodReads, we use the listed featured words for mystery and children’s books,⁴ and handcraft their replacements.

Privacy-Aware Text Rewriting (PATR) We re-implement Xu et al. (2019)’s adversarial rewriting model with Huggingface pretrained translators. We first translate English input to French⁵ mediated results and translate it back to English⁶.

A.2 Case Studies

Case Study 1 (Gender) In Table 5, besides the pronoun *her*, *adorable* is also a strong predictor of female gender (the word ‘*ributes*’ is a typo by the referrer). Whether the adjectives are gendered depends on context (e.g., “beautiful work” may not predict gender but “beautiful person” does). This is a difficult case for RB and WD to distinguish or to select the best replacements. ADV replaces the gendered but positive word *adorable* with a neutral but less positive word *third*. This reveals that while ADV substitutes a less biased word, it lacks the ability to maintain the high-level semantic meaning. PATR shows its advantage of paraphrasing due to the back-translation, however, it fails to identify biased words and debias them.

DEN and DEPEN successfully neutralize *adorable* → *commendable* or *praiseworthy* which express not only the same semantic meaning but also the same high-level sentiment. Noting that we don’t have any sentiment guidance or constraint,

³<https://writingcenter.unc.edu/tips-and-tools/gender-inclusive-language/>

⁴<https://www.vocabulary.com/lists/>

⁵<https://huggingface.co/Helsinki-NLP/opus-mt-en-fr>

⁶<https://huggingface.co/Helsinki-NLP/opus-mt-fr-en>

this advantage is achieved by grasping the core content and inferring the underlying attitude. DEN and DEPEN can correct the typo *ributes* with a plausible replacement (*work*).

Another interesting phenomenon is when DEPEN accidentally generates a gendered word (*Her*), it compensates by correcting this to a proper noun (essentially an ‘invented name’ *Her* → *Heragur*); the new word still plays the same grammatical role in the sentence (e.g., *Her* and *Heragur’s* are possessive pronouns with the same POS tag). This could perhaps be further improved by preventing the decoder from generating proper nouns at all, or otherwise by combining our decoding strategy with additional rules.

Case Study 2 (Gender) As shown in Table 6, *He*, *lover* and *basketball* are predictive of (male) gender. Although WD, ADV and DEN find a close replacement (*sports* for *basketball*), the sentences still predict the male gender (they fail to correct the pronoun *he*). While it replaces *lover* with a more neutral word (*enthusiast*), PART still generates *he* and *basketball*. PEN again rewrites the sentence in a way that differs drastically from the input. DEPEN neutralizes the highlighted parts with suitable replacements.

Case Study (Nationality) Table 7 shows a sentence in a reference letter written for a US student. We find that ‘extracurricular’ activities (both the word itself and the topic in general) tend to appear more in letters for US (and to some extent Indian) students compared to (e.g.) Chinese students; as such the word is detected as a predictor of nationality. From Table 7, although RB eliminates the indicator *extracurricular*, it causes ambiguity by simply deleting it. DEPEN replaces the indicator *extracurricular* with *social/cultural* which is not only semantically similar but also less predictive of nationality (note that the pronoun ‘her’ is not removed from this sentence as it is not a sensitive attribute in this experiment).

Case Study (GoodReads) Table 8 shows a sentence from a review of a children’s book, where models rewrite to hide genre information while maintaining content (especially the review sentiment). This example gives another illustration about why rule-based (RB) methods fail: *children* in this context does not refer to the genre but describes a specific character. Distinguishing such differences would demand a more nuanced rule-based

Sensitive Attr.	Rules
Gender	Replace he/she → they, his/him/her/hers → them/their, boy/girl → person Delete Mr., Ms., Miss, Mrs. Replace chairman/chairwoman → chair, actor/actress → actor, freshman → first-year student ...
Nationality	Delete country/city/nationality names, e.g., China/Chinese, America/American, India/Indian, Taiwan ... Category 1: Replace intellectual curiosity → ability Category 2: Replace solid foundation → understanding Category 3: Replace financial/finance aid/support/situation → support/situation Category 4: Replace senior project → project
Genre	Replace children/child/kid/boy/girl/daughter/son → reader, picture/children/fairy book/story → book/story Delete murder, mystery, crime, suspect, suspense, victim, killer, investigation ...

Table 4: Detailed replacement rules used in our rule-based baseline.

model, requiring significant handcrafting. DEPEND can overcome this problem by doing inference automatically.

A.3 Data Preprocessing

For the Reference Letter dataset, we first exclude invalid reference letters if the letter (1) is too short (less than 2 sentences), or (2) contains too many named entities (more than 90%, presumably due to OCR problems), or (3) is not written in English. For GoodReads dataset, we sample 3000 samples each from Children’s and Mystery’s genre.

A.4 Details of Model

A.4.1 Number of Parameters

In all experiments, we use BERT in the Detect stage, which has 110M parameters; we use BART as our base Seq2Seq model in the Perturb stage, which has 117M parameters. All classifiers are finetuned BERT.

A.4.2 Hyperparameters

We use 64 as the batch size for finetuning all BERT classifiers and use 8 as the batch size for the BART Seq2Seq model for finetuning or generation. We use AdamW⁷ as the optimizer with initial learning rate of 1e-4. The whole pipeline is implemented with PyTorch⁸, and all transformers are implemented based on the libraries of Hugging Face⁹.

In our Perturb stage, we tried several k ($k = 10, 20, 30$) during our implementation and we found that our results are not sensitive to the choice of k .

⁷<https://pytorch.org/docs/master/generated/torch.optim.AdamW.html>

⁸<https://pytorch.org/>

⁹<https://huggingface.co/>

A.5 Details of Datasets

We download the GoodReads book review dataset by genre from the official website¹⁰.

A.6 Details of Evaluation Metrics

`nltk.translate.bleu_score.corpus_bleu` from `nltk` package is used to calculate the BLEU4 scores.

We use the official repository¹¹ to calculate the Pseudo-Log-Likelihood scores of generated sentences.

¹⁰https://sites.google.com/eng.ucsd.edu/ucsdbookgraph/home#h.p_VCP_govwttn1

¹¹<https://github.com/awsmlabs/mlm-scoring.git>

Model	Re-generated
Original	Her desire for perfection, confidence levels, humility and excellent personal at ributes are adorable .
RB	Their desire for perfection, confidence levels, humility and excellent personal at ributes are adorable.
WD	Her desire for perfection, confidence levels, humility and excellent personal at ributes are adorable.
ADV	Her desire for perfection, confidence levels, humility and excellent personal at kuutes are third.
PATR	His desire for perfection, level of confidence, humility and excellent personal attributes are adorable.
DEN	Her desire for perfection, confidence levels, humility and excellent personal at work are commendable.
PEN	Forges says, desire, humility, humility and desire to be honest with perfect personal.
DEPEN	Heragur's desire for perfection, for perfection, confidence levels, humility and excellent personal at work are praiseworthy.

Table 5: Re-generated examples with gender as the sensitive attribute

Model	Re-generated
Original	Meanwhile he is not a keen lover of basketball , but also plays it with skills.
RB	Meanwhile they is not a keen lover of basketball, but also plays it with skills.
WD	Meanwhile he is not a keen learner of sports, but also plays it with skills.
ADV	Meanwhile he is not a keen learner of sports, but also plays it with skills.
PATR	Meanwhile he is not a basketball enthusiast, but also plays with skills.
DEN	Meanwhile he is not a keen lover of sports, but also plays it with skills.
PEN	The Duchess of Amida Costa Rica plays the World No. 3-rank seven in a row.
DEPEN	Meanwhile: PERSON-I-2189 is not a keen learner of sports, but also plays it with skills.

Table 6: Re-generated examples with gender as the sensitive attribute.

Model	Re-generated
Original	Apart from her studies she has also taken keen interest in extracurricular activities
RB	Apart from her studies she has also taken keen interest in activities .
WD	Apart from her classes she has also taken keen interest in co-curricular activities.
ADV	Apart from her studies she has also taken keen interest in extracurricular activities.
PATR	Apart from her studies, she also interested herself in extracurricular activities.
DEN	Apart from her coursework she has also taken keen interest in extra-curricular activities.
PEN	Samantha Smith is a great-and-groom at home to an amazing six-week-old.
DEPEN	Apart from her coursework she has also taken keen interest in social/cultural activities

Table 7: Re-generated examples with nationality as the sensitive attribute.

Model	Re-generated
Original	I didn't really get this one, although I liked the example of children dealing with a new sibling.
RB	I didn't really get this one, although I liked the example of readers dealing with a new sibling.
WD	I didn't really like this one, although I liked the story about kids dealing with a new sibling.
ADV	I didn't really get this one, although I liked the example of children dealing with a new sibling.
PATR	I didn't really get this one, though I liked the example of kids dealing with a new brother and sister.
DEN	I didn't really like this one, although I liked the idea of siblings dealing with a new sibling.
PEN	Young Wolf stories that deal with siblings siblings.
DEPEN	I don't really like this one, although I liked the story of characters dealing with a new sibling.

Table 8: Re-generated examples with genre as the sensitive attribute.