

Gender-Inclusive Grammatical Error Correction through Augmentation

Gunnar Lund, Kostiantyn Omelianchuk, Igor Samokhin

Grammarly

{gunnar.lund,kostiantyn.omelianchuk,igor.samokhin}@grammarly.com

Abstract

In this paper we show that GEC systems display gender bias related to the use of masculine and feminine terms and the gender-neutral singular *they*. We develop parallel datasets of texts with masculine and feminine terms, and singular *they*, and use them to quantify gender bias in three competitive GEC systems. We contribute a novel data augmentation technique for singular *they* leveraging linguistic insights about its distribution relative to plural *they*. We demonstrate that both this data augmentation technique and a refinement of a similar augmentation technique for masculine and feminine terms can generate training data that reduces bias in GEC systems, especially with respect to singular *they* while maintaining the same level of quality.

1 Introduction

Natural Language Processing (NLP) systems are well known to exhibit sensitivity to social characteristics, a sensitivity that may lead to harms for users interacting with these systems. In this work, we examine how NLP systems performing the task of Grammatical Error Correction (GEC) are sensitive to *gender* characteristics in English and how this sensitivity represents bias that harms users of these systems. We propose a new data augmentation technique to address bias related to singular *they*, and show how it can mitigate gender bias. We also show how this technique interacts with prior work on data augmentation methods to reduce bias.

The expression of gender in English is complex, and to focus our study, we identify discrete biases that GEC systems may exhibit. First, we build on other works on gender bias in NLP by examining how masculine and feminine terms impact the behavior of GEC systems. Discrepant behavior between sentences that contain one or the other is evidence of bias. Second, we examine the behavior of GEC systems with a gender-neutral pronominal

paradigm in English commonly called singular *they*. The linguistic properties of this paradigm introduce additional kinds of biases relative to masculine or feminine pronouns. Additionally, we make our evaluation datasets publicly available¹.

To prevent these adverse behaviors, we apply techniques to generate synthetic training data that address different aspects of these behaviors. First, we adopt Counterfactual Data Augmentation (CDA), which is used successfully to reduce bias in word embedding models (Lu et al., 2019; Maudslay et al., 2020), and apply it to the GEC case. Second, we introduce a new technique for generating training data with *they* pronouns that have unambiguously singular reference using insight from theoretical linguistics to target model bias with gender-neutral sentences. Because these techniques are data-oriented, the innovations should theoretically generalize beyond the GEC domain to other NLP tasks as well.

Our main contributions are:

- We introduce a novel technique for creating singular *they* data, leveraging specific linguistic features of this use of the pronoun. Additionally, we refine previous approaches to addressing discrepancies in model behavior between texts with masculine and feminine pronouns and show their application outside of the context of masked language models.
- We qualitatively and quantitatively measure biases in competitive GEC systems by comparing model performance on parallel test sets containing singular *they* pronouns, masculine terms, and feminine terms.
- We show how our data augmentation techniques, both in isolation and combination, mitigate these biases when used to create training

¹<https://github.com/grammarly/gender-inclusive-gec>

data for these GEC systems with a minimal impact to overall performance.

2 Background

2.1 Gender and Bias in NLP and GEC

2.1.1 Conceptual grounding

To orient this work, we highlight two recent calls-to-action regarding the study of bias and gender in NLP systems. First, following Devinney et al. (2022), we mean to be explicit about the conception of gender and gendered language we assume. In particular, we are concerned with gendered linguistic *content*, and not the gender of the authors or readers of that content. We recognize that the expression of gender in English is notional; the use of some nouns and pronouns is linked to particular gendered conceptual categories (McConnell-Ginet, 2013; Ackerman, 2019). Additionally, the *use* of language with gendered content represents one aspect of gender performativity which produces and reifies these gendered categories.

Second, following Blodgett et al. (2020), “bias” is an inherently normative concept. In the context of NLP systems, it must be understood in terms of the potential *harms* that those systems may cause and to *whom* those harms may be caused. Therefore, we directly focus on mitigating harms themselves as they relate to the GEC task and how users of these systems interact with them and may be affected by them. Unlike some studies of bias in upstream contexts like word embeddings, users interact with GEC systems directly; these users choose to incorporate these systems’ suggestions into their emails, essays, and tweets, and bias may impact anyone interacting with such text.

Further, because the GEC task is an inherently normative one on its own—these systems offer suggestions to *correct* a user’s text and are designed in accordance with preexisting normative notions of “correct” or “fluent” English—GEC systems necessarily also participate in the production of gendered categories. The norms assumed when constructing these systems and datasets in this regard may conflict with other norms about language use. For example, there are norms against the use of singular *they* in some language communities. Some English speakers do not accept singular *they* as a grammatical construction of English (Bjorkman, 2017, a.o.), and some prescriptive grammars advise against the use of singular *they* (c.f., Strunk and White, 1999). People who are non-binary and use *they* pronouns,

e.g., cannot refer to themselves “correctly” within these circumscribed norms of language use. As we discuss below, the operationalization of these norms in GEC systems may lead to harm. We adopt the view that GEC systems should reflect the most permissive distribution of singular *they*. This distribution is discussed further in section 3.2.

Our work has a notable limitation in that we do not investigate bias with respect to neopronouns like *ze* or *xe*. We leave extensions of CDA-like techniques to these pronouns for future work.

2.1.2 Two biases

We identify two areas where GEC systems can produce biased, and therefore potentially harmful, outcomes. Importantly, this is not an exhaustive account of potential biases GEC systems exhibit, but we think this is as good a starting point as any.

First, a GEC system can be implicitly biased if it consistently performs better on texts containing words of one gendered category over another. This is an allocative harm. If a GEC system performs worse on texts that are about people who use one pronoun or another, texts about those people may contain more grammatical errors, impacting their relative opportunity. For example, a user writing letters of recommendation may inadvertently include more grammatical errors in letters for individuals using masculine pronouns, as a system could perform worse on texts with masculine pronouns than feminine pronouns, and this could impact the relative reception those letters receive compared to similar letters with feminine pronouns.

Second, a GEC system can be explicitly biased if it offers corrections that reify harmful notions about particular gendered categories, including the reinforcement of stereotypes and misgendering or erasure of individuals referred to in the user’s text. This is a representational harm. This harm is explicitly called out by participants in a survey on harms of AI systems with respect to non-binary individuals (Dev et al., 2021). The examples below are representative of these kinds of corrections. The first is an instance of misgendering, replacing singular *they* with a masculine pronoun; the second is an instance of erasure, implying that *they* has a correct use only as a plural pronoun.

1. I asked Alex **their** phone number. -> I asked Alex **his** phone number.
2. They are **a linguist**. -> They are **linguists**.

We find evidence for both of these biases by analyzing the following GEC systems (table 1)

1. GECToR (Omelianchuk et al., 2020), sequence tagging approach, which was a state-of-the-art GEC model in 2020
2. Fine-tuned BART model (Lewis et al., 2020), which represents another popular and competitive approach - sequence to sequence
3. EditScorer (Sorokin, 2022), the recent ranker approach, that is the second-best result on BEA Shared Task 2019, as of April 2023².

Our quantitative analysis revealed that, for all three systems, there is a significant gap (from -6.2% to -9.5% F05 points) between the original and augmented with singular *they* examples versions of the BEA-dev subset, which we call bea-195. The details on how we built these datasets and evaluation approach is provided in sections 4.1 and 4.2. Detailed evaluation results are available in appendix table 7.

We hypothesize that these biases share a partial cause: an imbalance in the training data. If, e.g., the training data with masculine words is of a higher quality than that with feminine words, there may be a performance gap. In the case of unnecessary corrections of singular *they*, we hypothesize that the imbalance is caused by an extreme lack of singular *they* sentences relative to plural *they* sentences. In the remainder of this paper, we show that the introduction of synthetic data helps mitigate these biases.

2.2 Related work

2.2.1 Data augmentation

Data augmentation has been used in other NLP domains to mitigate gender bias, but most of these works focus on just the masculine and feminine gender categories in English and limit the application of these techniques to word embedding models. Zhao et al. (2018); Rudinger et al. (2018); Lu et al. (2019) show that coreference resolution systems are sensitive to masculine and feminine words in otherwise equivalent sentences. Lu et al. (2019) use what they call Counterfactual Data Augmentation (CDA) to reduce this sensitivity. In CDA, masculine pronouns are swapped for feminine ones

and vice versa. They also swap definitionally gendered common nouns like *actor* and *actress*. They set aside data where the swapping candidates are in a cluster with a proper name.

Maudslay et al. (2020) extend Lu et al. (2019) CDA and implement additional name swapping, where the gendered associations of names were determined using census data from the US Social Security Administration. They use this technique to minimize gendered differences in word embedding spaces as measured by WED (Bolukbasi et al., 2016). As in Lu et al., Maudslay et al. limit their method to masculine and feminine categories.

In addition, gender-neutral data augmentation methods have been proposed in concurrent works by Sun et al. (2021) and Vanmassenhove et al. (2021). These methods have different goals than ours and are designed to produce different kinds of data. We discuss the differences between these methods and our own in section 3.3.

To our knowledge, ours is the first work to use CDA techniques to reduce bias in GEC systems and the first to use singular *they* augmentation to inject synthetic training data to reduce bias with singular *they* sentences.

2.2.2 Singular *they* and NLP systems

Previous works investigating bias towards singular *they* sentences have generally focused on coreference resolution systems. Cao and Daumé III (2021) develop a dataset to evaluate these systems on naturalistic texts about individuals who identify as non-binary, where 35% of the pronouns are singular *they*. They report that the Stanford system is the highest scoring on this dataset with an F1 score of 34.3%. This same system reports a much higher F1 score of 60% on the CONLL 2012 test set.

Baumler and Rudinger (2022) compare coreference resolution system performance directly on singular *they* sentences compared to plural *they* sentences along the lines of the Winograd or Winogender schemata (Levesque et al., 2012; Rudinger et al., 2018; Zhao et al., 2018). They find across-the-board gaps in system performance between the two test sets.

Outside of coreference resolution, Dev et al. (2021) investigate biased representations with BERT in a masked word prediction task. They find that for masked pronouns, BERT has a high accuracy in the prediction of masculine and feminine pronouns, but accuracy considerably lowers for singular *they*.

²http://nlpprogress.com/english/grammatical_error_correction.html

System	bea-dev-full	bea-195			bea-556		
	F05	F05 orig	F05 st aug	diff	F05 orig	F05 mf aug	diff
GECToR (roberta-base)	54.57%	58.28%	48.74%	-9.54%	59.23%	58.96%	-0.27%
BART (seq2seq)	52.74%	56.36%	50.13%	-6.23%	58.61%	58.79%	0.18%
EditScorer (roberta-large)	58.92%	60.6%	54.16%	-6.44%	62.55%	61.59%	-0.96%

Table 1: Scores on BEA-dev subsets for strong GEC baselines.

3 Description of the data augmentation methods

We use two data augmentation methods. First, we follow Lu et al. (2019) and others in swapping out feminine words for masculine words and vice versa. Second, we propose a novel augmentation method for generating singular *they* data from sentences containing masculine and feminine pronouns. We treat singular *they* differently because language internal facts about English necessitate a separate treatment: unlike *he* and *she*, *they* has a second life as a plural pronoun.

3.1 Feminine/Masculine CDA (FM-CDA)

Consistent with masculine/feminine-term swapping methods in other works, we swap three kinds of nominal terms:

- Pronouns: Swap masculine pronouns for their feminine counterparts and vice versa. Ex: *him* → *her*. Because the masculine and feminine pronominal paradigms are partly syncretic—the feminine pronoun *her* can be accusative or possessive and map to *him* or *his*, respectively—token POS tags, generated by a proprietary POS tagger, were used to appropriately match terms to their case-same counterpart.
- Common nouns: Swap definitionally feminine common nouns for their masculine gendered counterparts and vice versa. Ex: *actor* → *actress*. The selection and mapping of these nouns were hand-curated by industry experts.
- Names: Swap first names that are usually associated with feminine terms for names usually associated with masculine terms and vice versa. We partnered with industry experts to curate dictionaries of masculine and feminine names. Because names don’t necessarily have gendered counterparts in the way that pronouns or common nouns do, an arbitrary mapping of names was created. Names occurring in both lists were excluded from swapping.

Unlike some previous work involving CDA on fully unsupervised tasks where the creation of a single counterpart sentence is sufficient, GEC training data consists of pairs of ungrammatical source text and grammatical target text. This introduces challenges similar to those that CDA faces for machine translation data consisting of parallel texts (Saunders and Byrne, 2020; Wang et al., 2022) Because, e.g., the POS tagger may perform differently on the two texts, especially given that the source text is ungrammatical, the swapping algorithm may produce inconsistent swaps if applied separately to the source and target texts. This inconsistency can introduce grammatical errors between the source and target texts and negatively impact model performance.

To avoid this, an additional algorithm ensures a consistent swap between the source and target text where possible. We first apply our algorithm to the grammatically corrected target text. Then we use an alignment algorithm to align the target and source texts and isolate the differing segments of the texts. For each differing segment, we determine if the number of tokens in the source and target segments is the same, and if not, we discard the data. Then we compare every token in the source and target segments; if the source word would have the same swap as the target, we replace the source word with its differently gendered counterpart. If neither word is swappable, we do nothing. If there is a mismatch in the swap between source word and target word, we discard the data point. For the evaluation set discussed below, we reintroduced this discarded data and manually edited the data to introduce singular *they* pronouns and ensure that the results are parallel.

3.2 Singular *they* CDA (St-CDA)

Singular *they* is an inherently referential phenomenon. As is evident in the name, it is distinguished from plural *they* because it refers to singular individuals. Further, theoretical and experimental linguistic works show that the overall distribution of singular *they* is conditioned by the nature of

the singular antecedent and discourse participants’ relation to the antecedent (Bjorkman, 2017; Ackerman, 2019; Moulton et al., 2020; Konnelly and Cowper, 2020; Han and Moulton, 2022). Speakers may be sensitive to linguistic vs. non-linguistic antecedents (Moulton et al., 2020), specificity and definiteness of the antecedent (Bjorkman, 2017; Konnelly and Cowper, 2020), the discourse participant’s knowledge of the referent’s gender identity (Bjorkman, 2017; Ackerman, 2019; Konnelly and Cowper, 2020), and the association of a lexical item or name with a particular gender category (Bjorkman, 2017; Ackerman, 2019; Moulton et al., 2020). In the case of the broadest distribution, singular *they* is used in the same ways that masculine and feminine pronouns are used—the referent’s preference largely dictates the choice of pronoun—but *they* may additionally be used when the referent’s preference is not known (Konnelly and Cowper, 2020).

Differently than feminine and masculine pronouns, we hypothesize that adverse model behavior with singular *they* is at least partially caused by its infrequency relative to plural *they*. Therefore, it is not enough to simply create data that has *they* pronouns; it must also be evidently singular as well. We leverage these linguistic insights about antecedenthood to identify contexts where swapping will result in unambiguous cases of singular *they* by identifying singular antecedents of the pronouns in the text.

We implement this by using HuggingFace’s Neuralcoref coreference resolution system³ built on top of SpaCy⁴. For a given coreference cluster with a masculine or feminine pronoun, we look at the coreferring expressions in the cluster. If we find a singular one, we perform the swap. We consider a coreferring expression singular if it is:

1. A singular common or proper noun.
2. A singular possessum (e.g., *his foot*).

In addition, *they* has different verbal agreement paradigms than *he* or *she*. We resolve this by using SpaCy’s dependency parser to identify agreeing verbs with the swapped pronouns. We then use the pyInflect package⁵ to select the verbal inflection consistent with subject agreement with *they*.

³<https://github.com/huggingface/neuralcoref>

⁴Neuralcoref works with SpaCy v2.1 (<https://v2.spacy.io/>).

⁵<https://github.com/bjascob/pyInflect>

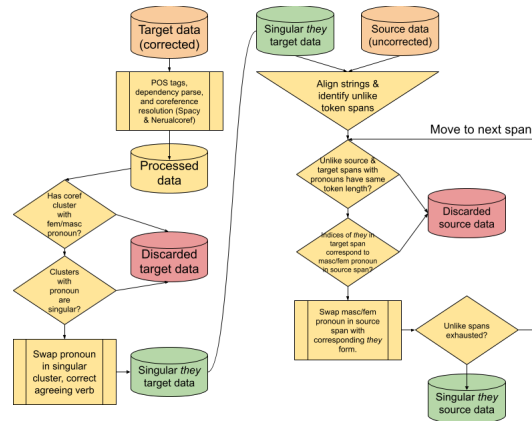


Figure 1: A graphical illustration of the St-CDA swapping process.

Finally, we use POS information to disambiguate syncretic forms of *her*. For reflexive pronouns, we swap in the form *themselves* and not *themselves* as this form is much less common in the preexisting training data and is more likely to lead to a singular interpretation of *they* pronouns.

As in the case of FM-CDA, we face the potential for inconsistency if this technique is used separately for both source and target data. We use the same algorithm as FM-CDA to perform safe swaps with an additional check on the verbs that were corrected to agree with *they* in the target swapped data.

3.3 Differences between St-CDA and other approaches

To our knowledge, there are two similar approaches to generating singular *they* data in (Sun et al., 2021) and (Vanmassenhove et al., 2021). Their approaches differ from ours in two crucial ways. First, their techniques are meant to create wholly gender-neutral texts, swapping out instances of definitionally gendered noun phrases like “fireman” for “firefighter”. We do not perform these swaps because *they* pronouns may corefer with definitionally gendered words, and this data, in particular, is likely to be rare in most corpora. As we see qualitatively, the baseline model we test seems to be particularly likely to “correct” singular *they* sentences unnecessarily when there is a definitionally gendered word in the sentence.

Second, and most crucially, we hypothesize that the performance gap with singular *they* sentences is due to the relative glut of plural *they* data compared to singular *they* data. As such, we seek to add *they* data that unambiguously has singular reference. Their techniques, by contrast, may result in

data where *they* may have a primarily plural interpretation. By targeting contexts where *they* is more likely to be interpreted singularly, we believe St-CDA produces data that will have a higher positive impact and have fewer adverse effects on model quality.

Ultimately, Sun et al. (2021) and Vanmassenhove et al. (2021) have different goals than we do, and this informs the differences in our techniques. Both works envision their technique to be used at runtime in machine translation tasks to, e.g., ensure translations from languages with grammatical gender result in gender neutral translations in English. While they speculate that their techniques can be used to create augmented training data, as we do in this paper, they do not specify what issues they intend this augmented training data to address. By contrast, we seek to counteract a particular imbalance between plural and singular *they* sentences.

4 Experiments with GEC

4.1 Description of datasets

For training data, we chose a large Lang-8 Corpus of Learner English (Mizumoto et al., 2011), and more specifically, its “cleaned” version cLang-8 (Rothe et al., 2021) which contains over 2 million corrected English sentences. The downside is the noisiness of the data (even in the “cleaned” version) and the lack of consistency in annotations. There are few other GEC datasets of comparable size (Bryant et al., 2022). Naturally, not every sentence contains personal pronouns, so only a subset of the dataset is suitable for data augmentation. The size of cLang-8 allowed us to produce about 63 thousand sentences with singular-they augmentation and 254 thousand gender-swapped sentences, which is enough for fine-tuning purposes. In our further experiments, we used only a random sample of 50 thousand sentences from each augmented version of the data to make sure that the results were not impacted by a difference in data size.

4.2 Evaluation approach

4.2.1 Description of the evaluation procedure

There is no evaluation set which would specifically contain multiple uses of the singular *they*, so we need to apply data augmentation here as well. To do this, we use the dev part of the BEA-2019 shared task (Bryant et al., 2019) since it is one of the standard evaluation sets for GEC. Of 4384 sentences in the BEA dataset, 195 singular *they*

sentences were created by replacing the pronouns “he” and “she” with singular “they.” To do this, we applied the CDA-st algorithm described above. The data discarded by the alignment algorithm was also collected and manually revised where possible (sentences where, e.g., a pronoun was inserted or changed from one gendered pronoun to another were either eliminated or revised to eliminate the error). Finally, the entire dataset was manually reviewed to ensure consistency between the original data and the augmented data.

To find the difference in GEC performance on sentences with and without the singular “they,” we evaluate on the subset of 195 sentences before augmentation, “BEA-195-orig”, and on the 195 augmented sentences, “BEA-195-st-aug”. The dataset size limits the conclusions we can make about the GEC model’s performance in general, but the differences between scores obtained by the GEC models on these two subsets are statistically significant.

We repeat this procedure for experiments involving masculine and feminine swapping. In this case, our augmentation produced subsets of 556 sentences: “BEA-556-orig” and “BEA-556-mf-aug”.

4.2.2 Error distribution analysis

To ensure that our augmentation did not affect edits and shift the error distribution, we conducted a qualitative analysis of m2 files produced by Errant tool on parallel sentences of the original and augmented versions of “195” and “556” evalsets. As shown in the edit type distribution (appendix section B.1), there are only minor differences in the number of edits (less than 1% of edits affected). It can be explained by the fact that sometimes Errant might represent similar edits by single or multiple edits, like in the following example 2.

The error type distribution for both subsets is available in appendix section B.2.

4.2.3 Questions to answer with evaluation

Running evaluation on these datasets, we are interested mainly in answering two questions:

- Is the state-of-the-art GEC model, which was not trained specifically with singular “they” or gender-swapped data, producing worse corrections on the augmented evaluation dataset?
- If the corrections are worse, can we shrink or remove the gap in performance by fine-tuning the model on the augmented training data?

Data source	Sentence	Edits
original bea-dev	I love this game because my favourite sport man belong to this game .	sport man belong to this => sportsman plays
mf aug bea-dev	I love this game because my favourite sport woman belong to this game .	sport woman => sportswoman belong to this => plays

Table 2: An example of a sentence with a different number of edits in an m2 file depending on data augmentation.

4.3 Description of models

For experiments, we use GECToR (Omelianchuk et al., 2021) - a state-of-the-art GEC model based on the efficient sequence tagging approach to corrections. Instead of producing a new error-free sentence, GECToR predicts a sequence of tags denoting operations: “keep,” “remove,” “insert_X,” or “append_X.” The corrected text is reconstructed from the original sentence and the tags. Sequence tagging is computationally cheaper than autoregressive approaches, which makes GECToR up to ten times faster than sequence-to-sequence models. At the same time, GECToR set the state-of-the-art at the time of publication.

GECToR is trained and fine-tuned in several stages, starting from the pre-trained language model such as RoBERTa (Liu et al., 2019). One can also start from the fine-tuned GECToR checkpoint (available on GitHub) and fine-tune it further on the data specifically tailored to the task at hand. However, it may lead the catastrophic forgetting issue, and the overall performance of the model on the general GEC test sets may deteriorate.

4.4 Experiment approach

We select GECToR for our fine-tuning experiments due to it being a competitive GEC system and having code that is publicly available. We use weights of the pre-trained GECToR (with RoBERTa-base encoder) model as initialization and fine-tune it for 5 epochs on the following data:

1. Original clang8 sentences (2.2m sentences)
2. Mix of original and augmented clang8 sentences of one type (2.2m + 50k sentences, either singular-they or gender-swapped)
3. Mix of original and augmented clang8 sentences of both types (2.2m + 100k sentences, 50k for both singular-they and gender-swapped)

We fine-tune the model for 5 epochs with early stopping after 3 epochs and 1 cold epoch. For

each training data configuration, we run training 10 times with different random seeds and report the average across all run results. The full list of hyperparameters for fine-tuning can be found in Appendix B.

Because the baseline GECToR model is already strong enough (it was a SOTA model in 2020) and clang-8 is high-quality data produced by another strong GEC system gT5 xxl (Rothe et al., 2021), the fine-tuning does not lead to substantial quality degradation. As shown in table 3, the differences in F0.5 scores are statistically insignificant.

#	Used clang data			bea-dev (full)
	Orig	MF	ST	F05 orig
0	no	no	no	54.58%
1	yes	no	no	54.61% ± 0.41%
2	yes	no	yes	54.52% ± 0.48%
3	yes	yes	no	54.44% ± 0.58%
4	yes	yes	yes	54.63% ± 0.56%

Table 3: F0.5 on BEA-dev-full for GECToR fine-tuning experiments. For new experiments, average over all seeds ± 2 s.d. is shown.

To evaluate the impact of adding the augmented data to the training dataset, we used an original subset of BEA dev and its augmented manually reviewed versions (described above). The results are shown in table 4 and table 5.

4.4.1 Experiment with singular-they augmentation

We can see that for the baseline model, the gap in F0.5 between the original and augmented (singular-they) version of BEA dev subset is quite significant -9.54%. Fine-tuning on clang8 data led to the shrinking of the gap to -5.86%. We think that this decrease illustrates not an improvement in gender bias, but rather a change in the baseline value due to a shift in precision/recall after fine-tuning. For a more fair comparison, we focused on analyzing the difference between the fine-tuned model on a

#	Used clang data			bea-dev 556		
	Orig	MF	ST	F05 orig	F05 mf_aug	Delta
0	no	no	no	59.23%	58.96%	-0.27%
1	yes	no	no	57.79% ± 0.82%	57.08% ± 1.06%	-0.71%
2	yes	no	yes	57.58% ± 1.12%	57.01% ± 1.2%	-0.57%
3	yes	yes	no	57.88% ± 0.7%	57.33% ± 0.78%	-0.55%
4	yes	yes	yes	58.03% ± 0.76%	57.5% ± 1.04%	-0.53%

Table 4: F0.5 on BEA-dev-556 for GECToR fine-tuning experiments. For new experiments, average over all 10 seeds ± 2 s.d. is shown.

#	Used clang data			bea-dev 195		
	Orig	MF	ST	F05 orig	F05 st_aug	Delta
0	no	no	no	58.28%	48.74%	-9.54%
1	yes	no	no	56.33% ± 2.1%	50.47% ± 1.62%	-5.86%
2	yes	no	yes	55.71% ± 1.22%	54.31% ± 1.62%	-1.4%
3	yes	yes	no	55.77% ± 1.04%	50.33% ± 1.12%	-5.44%
4	yes	yes	yes	56.33% ± 1.48%	54.86% ± 1.46%	-1.47%

Table 5: F0.5 on BEA-dev-195 for GECToR fine-tuning experiments. For new experiments, average over all 10 seeds ± 2 s.d. is shown.

combination of original and augmented data from clang8 (systems 2,3,4) and a model fine-tuned only on original clang data (system 1) (table 5).

We got an improvement in the F0.5 gap for systems 2 and 4 (from -5.86% to -1.4% and -1.47% correspondingly). This reduction is driven by the improvement on the augmented version of bea-dev-195 subset (F0.5 +3.84% and +4.39%) without any (0% for system 4) or with insignificant degradation in quality on the original bea-dev-195 subset (-0.62% for system 3).

We also qualitatively examine the corrections to determine whether explicit instances of bias are reduced through data augmentation (table 6). A linguist manually reviewed model predictions on bea-dev-195-st-aug for systems 1-4 and annotated predictions exhibiting explicit bias, which was defined as pluralization of a referent coreferring with singular *they* or the replacement of singular *they* with a gendered pronoun, or the replacement of *themselves* with *themselves*. System 4 shows the greatest improvement with 7 cases over the baseline of 32. Examples of explicit bias are in Appendix A.

4.4.2 Experiment with feminine/masculine augmentation

For feminine/masculine augmentation, the initial gap between the original subset of BEA (556 sentences) and the augmented version is much smaller -0.71%. Fine-tuning on original and femi-

#	Used clang data			bea-dev-195-st-aug	
	Orig	MF	ST	#	# w/o refl
1	yes	no	no	32	29
2	yes	no	yes	8	7
3	yes	yes	no	34	30
4	yes	yes	yes	7	4

Table 6: Number of sentences displaying explicit bias in bea-dev-195-st-aug. First column is total sentences found to have explicit bias, second is that count minus cases of "themselves">"themselves".

nine/masculine augmentation data (system 3) very slightly reduces this difference only to -0.55%. It's interesting that even singular-they augmentation, without any other gender-swapping, seems to provide a very similar result (difference of -0.57%). However, given the size of confidence intervals, we cannot say that any of our experiments had a significant impact on the gap.

4.4.3 Experiment with both augmentations

Finally, we tried to apply both kinds of augmentation - singular-they and feminine/masculine CDA. The resulting model (system 4) is producing very similar results in terms of gap difference for both BEA subsets that we used: -1.47% on bea-195 (system 2 gap is -1.4%) and -0.53% on bea-556 (system 3 gap is -0.55%), which is showing that

multiple biases might be handled with such a single fine-tuning approach at once. It also seems that augmented training data of two kinds does not interfere with any one evaluation but also does not provide additional benefits from this data interaction.

We believe that there are many other potential possibilities to incorporate augmented data into different stages of the training or change the proportion or the absolute number of original and augmented sentences in training data that might lead to even better improvement with little to no quality degradation. We would like to explore some of them in future work.

5 Conclusion

In this work, we developed a novel technique for data augmentation with sentences containing *they* that has an unambiguous singular reference and applied it to the GEC case. We used this technique to help develop a dataset of singular *they* data to parallel data in the BEA shared task dataset that has masculine and feminine pronouns, and with this, we show that GEC systems display bias in their treatment of singular *they* sentences compared to sentences with masculine or feminine pronouns. Additionally, we demonstrated that this technique could be used to reduce bias in GEC systems by fine-tuning the GEC system on the generated synthetic training data.

Because this technique is data-oriented, we believe that it has wider applications, and other NLP systems that display degraded performance with respect to singular *they* may benefit from being trained on data created through this technique.

Limitations

As noted, this work is limited in that it does not address neopronouns. We speculate that the augmentation techniques deployed in this work may extend to these pronouns as well, we recognize that they do not have the same linguistic reality as *he/shel/they* pronouns. Neopronouns may be similar to singular *they* in being relatively infrequent in a naturalistic corpus, but they are also different in that they don't overlap with a frequent morphologically-identical paradigm like plural *they*.

Additionally, the singular *they* augmentation technique we propose is specific to English and distributional facts about English pronouns. For one, English singular *they* morphologically overlaps with a plural pronoun, which is the primary

motivation for using coreference information to identify contexts where *they* would have a primarily singular interpretation. This is often not the case for other languages, as in Swedish where the gender-neutral *hen* is functionally similar to singular *they* but morphologically and distributionally dissimilar in that it does not overlap with a plural pronoun (Gustafsson Sendén et al., 2015).

Ethics Statement

Dataset risks

We do not anticipate any risks in releasing the evaluation dataset. This dataset was constructed through the modification of a publicly available dataset commonly used in the evaluation of GEC systems, the dev set of the BEA-2019 shared task (Bryant et al., 2019). These modifications involve the change of gendered words and agreeing verbs to create parallel data across masculine, feminine, and singular *they* pronouns with the goal of evaluating bias in GEC systems. By enabling researchers to measure bias in this way, we believe that the release of this dataset will aid further study in reducing bias in these systems by providing a benchmark.

Risks of describing data augmentation techniques

We caution that the singular *they* data augmentation technique used in this paper was not designed to generate text that surfaces directly to users. There may be risks to deploying data augmentation techniques at runtime as these techniques are designed to modify gender identity terms; depending on the context of deployment, users may be harmed by such modifications if they result in misgendering or erasure. On the other hand, as we show in this work, use of these techniques to generate training data can reduce bias, and we believe that in this way, the description of this technique will aid in reducing bias in NLP systems.

Acknowledgements

This research was supported by Grammarly. We thank our colleagues Leonardo Neves, Jade Razzaghi, Yichen Mo, Serhii Yavnyi, and Knar Hovakimyan for their brilliant insights and suggestions over the course of this work. We would also like to thank 3 anonymous reviewers for their helpful comments.

References

- Lauren Ackerman. 2019. [Syntactic and cognitive issues in investigating gendered coreference](#). *Glossa: a journal of general linguistics*, 4(1). Number: 1 Publisher: Open Library of Humanities.
- Connor Baumler and Rachel Rudinger. 2022. [Recognition of They/Them as Singular Personal Pronouns in Coreference Resolution](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3426–3432, Seattle, United States. Association for Computational Linguistics.
- Bronwyn M. Bjorkman. 2017. [Singular they and the syntactic representation of gender in English](#). *Glossa: a journal of general linguistics*, 2(1). Number: 1 Publisher: Open Library of Humanities.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(Technology\) is Power: A Critical Survey of "Bias" in NLP](#).
- Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. [Man is to Computer Programmer as Woman is to Home-maker? Debiasing Word Embeddings](#). Number: arXiv:1607.06520 arXiv:1607.06520 [cs, stat].
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019. [The BEA-2019 Shared Task on Grammatical Error Correction](#). In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75, Florence, Italy. Association for Computational Linguistics.
- Christopher Bryant, Zheng Yuan, Muhammad Reza Qorib, Hannan Cao, Hwee Tou Ng, and Ted Briscoe. 2022. [Grammatical Error Correction: A Survey of the State of the Art](#). ArXiv:2211.05166 [cs].
- Yang Trista Cao and Hal Daumé III. 2021. [Toward Gender-Inclusive Coreference Resolution: An Analysis of Gender and Bias Throughout the Machine Learning Lifecycle*](#). *Computational Linguistics*, 47(3):615–661. Place: Cambridge, MA Publisher: MIT Press.
- Sunipa Dev, Masoud Monajatipoor, Anaelia Ovalle, Arjun Subramonian, J. M. Phillips, and Kai Wei Chang. 2021. [Harms of Gender Exclusivity and Challenges in Non-Binary Representation in Language Technologies](#). In *EMNLP*.
- Hannah Devinney, Jenny Björklund, and Henrik Björklund. 2022. [Theories of "Gender" in NLP Bias Research](#). In *2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, pages 2083–2102, New York, NY, USA. Association for Computing Machinery.
- Marie Gustafsson Sendén, Emma A. Bäck, and Anna Lindqvist. 2015. [Introducing a gender-neutral pronoun in a natural gender language: the influence of time on attitudes and behavior](#). *Frontiers in Psychology*, 6.
- Chung-hye Han and Keir Moulton. 2022. [Processing bound-variable singular they](#). *Canadian Journal of Linguistics/Revue canadienne de linguistique*, 67(3):267–301. Publisher: Cambridge University Press.
- Lex Konnelly and Elizabeth Cowper. 2020. [Gender diversity and morphosyntax: An account of singular they](#). *Glossa: a journal of general linguistics*, 5(1). Number: 1 Publisher: Open Library of Humanities.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The Winograd schema challenge](#). In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning, KR'12*, pages 552–561, Rome, Italy. AAAI Press.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). ArXiv:1907.11692 [cs].
- Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2019. [Gender Bias in Neural Natural Language Processing](#). arXiv:1807.11714 [cs]. ArXiv: 1807.11714.
- Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. 2020. [It's All in the Name: Mitigating Gender Bias with Name-Based Counterfactual Data Substitution](#). arXiv:1909.00871 [cs]. ArXiv: 1909.00871.
- Sally McConnell-Ginet. 2013. ["Gender and its relation to sex: The myth of 'natural' gender"](#). In *"Gender and its relation to sex: The myth of 'natural' gender"*, pages 3–38. De Gruyter Mouton.
- Tomoya Mizumoto, Mamoru Komachi, Masaaki Nagata, and Yuji Matsumoto. 2011. [Mining Revision Log of Language Learning SNS for Automated Japanese Error Correction of Second Language Learners](#). In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 147–155, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.

- Keir Moulton, Chung-hye Han, Trevor Block, Holly Gendron, and Sander Nederveen. 2020. [Singular they in context](#). *Glossa: a journal of general linguistics*, 5(1). Number: 1 Publisher: Open Library of Humanities.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzhanyski. 2020. [GECToR – Grammatical Error Correction: Tag, Not Rewrite](#). ArXiv:2005.12592 [cs].
- Kostiantyn Omelianchuk, Vipul Raheja, and Oleksandr Skurzhanyski. 2021. [Text Simplification by Tagging](#). In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 11–25, Online. Association for Computational Linguistics.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2021. [A Simple Recipe for Multilingual Grammatical Error Correction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 702–707, Online. Association for Computational Linguistics.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender Bias in Coreference Resolution](#). arXiv:1804.09301 [cs]. ArXiv: 1804.09301.
- Danielle Saunders and Bill Byrne. 2020. [Reducing gender bias in neural machine translation as a domain adaptation problem](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7724–7736, Online. Association for Computational Linguistics.
- Alexey Sorokin. 2022. [Improved grammatical error correction by ranking elementary edits](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11416–11429, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- William Strunk and E. B. White. 1999. *The elements of style*, 4th ed edition. Allyn and Bacon, Boston.
- Tony Sun, Kellie Webster, Apu Shah, William Yang Wang, and Melvin Johnson. 2021. [They, Them, Theirs: Rewriting with Gender-Neutral English](#). ArXiv:2102.06788 [cs].
- Eva Vanmassenhove, Chris Emmery, and Dimitar Shterionov. 2021. [NeuTral Rewriter: A Rule-Based and Neural Approach to Automatic Rewriting into Gender Neutral Alternatives](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8940–8948, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jun Wang, Benjamin Rubinstein, and Trevor Cohn. 2022. [Measuring and Mitigating Name Biases in Neural Machine Translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2576–2590, Dublin, Ireland. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods](#). arXiv:1804.06876 [cs]. ArXiv: 1804.06876.

A Examples of explicit bias

Here we present examples of how the systems investigated here exhibit explicit bias. These are indicated in **bold text**. The list numbers correspond to the systems indicated in Table 3.

Source text:

I said that because when I was eleven my best friend had been operated on **their** back and before the operation he came with me and every day I had to wait for them because they spendt a lot of time in the shower cleaning their long hais , I hated that !

0. I said that because when I was eleven my best friend had been operated on and before the operation he came with me and every day I had to wait for them because they spendt a lot of time in the shower cleaning their long hair . I hated that !
1. I said that because when I was eleven my best friend had been operated on **his** back and before the operation he came with me and every day I had to wait for them because they spendt a lot of time in the shower cleaning their long hair . I hated that !
2. I said that because when I was eleven my best friend had been operated on their back and before the operation he came with me and every day I had to wait for them because they spendt a lot of time in the shower cleaning their long hair . I hated that !
3. I said that because when I was eleven my best friend had been operated on

their back and before the operation he came with me and every day I had to wait for them because they spent a lot of time in the shower cleaning their long hair . I hated that !

4. I said that because when I was eleven my best friend had been operated on their back and before the operation he came with me and every day I had to wait for them because they spent a lot of time in the shower cleaning their long hair . I hated that !

Source text:

To take my friend , Taylor , as an example , their homeland were over - developed and become an arid place when they were **a child** .

0. To take my friend , Taylor , as an example , their homeland was over - developed and became an arid place when they were **children** .
1. To take my friend , Taylor , as an example , their homeland was over - developed and became an arid place when they were **children** .
2. To take my friend , Taylor , as an example , their homeland was over - developed and became an arid place when they were a child .
3. To take my friend , Taylor , as an example , their homeland was over - developed and became an arid place when they were **children** .
4. To take my friend , Taylor , as an example , their homeland was over - developed and became an arid place when they were a child .

Source text:

My father has the same program in **their** computer and is able to buy on line tickets w / out leaving the house and moreover w / out picking up the phone .

0. My father has the same program on **his** computer and is able to buy online tickets w / out leaving the house and moreover w / out picking up the phone .

1. My father has the same program on **his** computer and is able to buy online tickets w / out leaving the house and moreover w / out picking up the phone .
2. My father has the same program on their computer and is able to buy online tickets without leaving the house and moreover without picking up the phone .
3. My father has the same program on **his** computer and is able to buy online tickets w / out leaving the house and moreover w / out picking up the phone .
4. My father has the same program on their computer and is able to buy online tickets without leaving the house and moreover without picking up the phone .

B Error distribution on BEA-dev subsets

B.1 Edit type distribution on "195" and "556" subsets of bea_dev

	bea_dev_195				bea_dev_556			
	orig		st_aug		orig		mf_aug	
Edit type	# edits	% edits	# edits	% edits	# edits	% edits	# edits	% edits
R (replacement)	374	66.1%	371	66.0%	814	64.1%	814	64.1%
M (missing)	156	27.6%	157	27.9%	358	28.2%	358	28.2%
U (unnecessary)	36	6.4%	34	6.0%	97	7.6%	96	7.6%
Total	566	100%	562	100%	1269	100%	1268	100%

B.2 Error categories distribution on "195" and "556" subsets of bea_dev

	bea_dev_195				bea_dev_556			
	orig		st_aug		orig		mf_aug	
Error category	# edits	% edits	# edits	% edits	# edits	% edits	# edits	% edits
PUNCT	150	26.5%	150	26.7%	314	24.7%	314	24.8%
VERB:TENSE	68	12.0%	66	11.7%	139	11.0%	139	11.0%
OTHER	45	8.0%	48	8.5%	121	9.5%	122	9.6%
PREP	46	8.1%	46	8.2%	115	9.1%	114	9.0%
DET	35	6.2%	34	6.0%	92	7.2%	89	7.0%
ORTH	44	7.8%	44	7.8%	84	6.6%	84	6.6%
SPELL	42	7.4%	41	7.3%	73	5.8%	74	5.8%
VERB	31	5.5%	30	5.3%	67	5.3%	67	5.3%
VERB:FORM	14	2.5%	15	2.7%	36	2.8%	37	2.9%
PRON	10	1.8%	10	1.8%	35	2.8%	36	2.8%
NOUN	17	3.0%	17	3.0%	34	2.7%	35	2.8%
NOUN:NUM	6	1.1%	6	1.1%	33	2.6%	32	2.5%
VERB:SVA	12	2.1%	11	2.0%	25	2.0%	25	2.0%
MORPH	8	1.4%	7	1.2%	24	1.9%	23	1.8%
ADV	12	2.1%	12	2.1%	16	1.3%	17	1.3%
ADJ	10	1.8%	9	1.6%	17	1.3%	16	1.3%
WO	5	0.9%	5	0.9%	12	0.9%	12	0.9%
NOUN:POSS	3	0.5%	3	0.5%	10	0.8%	10	0.8%
PART	5	0.9%	5	0.9%	8	0.6%	8	0.6%
CONTR	1	0.2%	1	0.2%	6	0.5%	6	0.5%
CONJ	2	0.4%	2	0.4%	6	0.5%	6	0.5%
VERB:INFL	0	0.0%	0	0.0%	2	0.2%	2	0.2%
Total	566	100%	562	100%	1269	100%	1268	100%

C Hyperparameter values for the fine-tuning of GECToR

Hyperparameter name	Hyperparameter value
batch_size	32
accumulation_size	4
n_epoch	5
patience	3
max_len	5
lr	1e-05
cold_steps_count	1
cold_lr	0.001
tp_prob	1
tn_prob	1
updates_per_epoch	10000
special_tokens_fix	1
transformer_model	roberta-base
Pretrained model	
Inference tweaks:	
minimum error probability	0.5
Inference tweaks:	
confidence	0.2

D Hyperparameter values for the fine-tuning of BART

Hyperparameter name	Hyperparameter value
base_model	BART-Large
src_max_length	80
tgt_max_length	85
beam	2
max_update	16000
loss_criterion	label_smoothed_cross_entropy
optimizer	Adam
weight_decay	0.0
adam_betas	(0.9, 0.98)
adam_eps	1e-06
lr	3e-05

E Full results of evaluation

System	bea-dev-full			bea-195-orig			bea-195-st-aug			bea-556-orig			bea-556-mf-aug		
	P	R	F05	P	R	F05	P	R	F05	P	R	F05	P	R	F05
GECToR (roberta-base)	64.05%	34.28%	54.57%	70.11%	34.81%	58.28%	56.33%	31.67%	48.74%	69.46%	37.27%	59.23%	68.84%	37.46%	58.96%
BART (seq2seq)	57.46%	39.7%	52.74%	61.32%	42.58%	56.63%	53.59%	39.86%	50.13%	62.73%	46.41%	58.61%	63.05%	46.29%	58.79%
EditScorer (roberta-large)	70.29%	35.77%	58.92%	73.98%	35.16%	60.6%	63.76%	33.81%	54.16%	75%	37.59%	62.55%	73.42%	37.46%	61.59%

Table 7: Scores on BEA-dev subsets for strong GEC baselines.

#	Used clang data			bea-dev (full)		
	Orig	MF	ST	Precision	Recall	F05 orig
0	no	no	no	64.05%	34.28%	54.58%
1	yes	no	no	62.29% \pm 1.3%	36.6% \pm 1.58%	54.61% \pm 0.41%
2	yes	no	yes	62.19% \pm 1.12%	36.35% \pm 0.98%	54.52% \pm 0.48%
3	yes	yes	no	62.38% \pm 1.1%	36.25% \pm 1.34%	54.44% \pm 0.58%
4	yes	yes	yes	62.41% \pm 0.72%	36.46% \pm 0.82%	54.63% \pm 0.56%

Table 8: F0.5 on BEA-dev-full for GECToR fine-tuning experiments. For new experiments, average over all seeds \pm 2 s.d. is shown.