

Inherent Biases in Reference-based Evaluation for Grammatical Error Correction and Text Simplification

Leshem Choshen¹ and Omri Abend²

¹School of Computer Science and Engineering, ² Department of Cognitive Sciences
The Hebrew University of Jerusalem

leshem.choshen@mail.huji.ac.il, oabend@cs.huji.ac.il

Abstract

The prevalent use of too few references for evaluating text-to-text generation is known to bias estimates of their quality (henceforth, *low coverage bias* or LCB). This paper shows that overcoming LCB in Grammatical Error Correction (GEC) evaluation cannot be attained by re-scaling or by increasing the number of references in any feasible range, contrary to previous suggestions. This is due to the long-tailed distribution of valid corrections for a sentence. Concretely, we show that LCB incentivizes GEC systems to avoid correcting even when they can generate a valid correction. Consequently, existing systems obtain comparable or superior performance compared to humans, by making few but targeted changes to the input. Similar effects on Text Simplification further support our claims.

1 Introduction

Evaluation in monolingual translation (Xu et al., 2015; Mani, 2009) and in particular in GEC (Tetreault and Chodorow, 2008; Madnani et al., 2011; Felice and Briscoe, 2015; Bryant and Ng, 2015; Napoles et al., 2015) has gained notoriety for its difficulty, due in part to the heterogeneity and size of the space of valid corrections (Chodorow et al., 2012; Dreyer and Marcu, 2012). Reference-based evaluation measures (RBM) are the common practice in GEC, including the standard M^2 (Dahlmeier and Ng, 2012), GLEU (Napoles et al., 2015) and I-measure (Felice and Briscoe, 2015).

The Low Coverage Bias (LCB) was previously discussed by Bryant and Ng (2015), who showed that inter-annotator agreement in producing ref-

erences is low, and concluded that RBMs underestimate the performance of GEC systems. To address this, they proposed a new measure, Ratio Scoring, which re-scales M^2 by the inter-annotator agreement (i.e., the score of a human corrector), interpreted as an upper bound.

We claim that the LCB has more far-reaching implications than previously discussed. First, while we agree with Bryant and Ng (2015) that a human correction should receive a perfect score, we show that LCB does not merely scale system performance by a constant factor, but rather that some correction policies are less prone to be biased against. Concretely, we show that by only correcting closed class errors, where few possible corrections are valid, systems can outperform humans. Indeed, in Section 2.3 we show that some existing systems outperform humans on M^2 and GLEU, while only applying few changes to the source.

We thus argue that the development of GEC systems against low coverage RBMs disincentivizes systems from making changes to the source in cases where there are plentiful valid corrections (open class errors), as necessarily only some of them are covered by the reference set. To support our claim we show that (1) existing GEC systems under-correct, often performing an order of magnitude less corrections than a human does (§3.2); (2) increasing the number of references alleviates under-correction (§3.3); and (3) under-correction is more pronounced in error types that are more varied in their valid corrections (§3.4).

A different approach for addressing LCB was taken by (Bryant and Ng, 2015; Sakaguchi et al., 2016), who propose to increase the number of references (henceforth, M). In Section 2 we estimate the distribution of corrections per sentence, and find that increasing M is unlikely to overcome LCB, due to the vast number of valid corrections

for a sentence and their long-tailed distribution. Indeed, even short sentences have over 1000 valid corrections on average. Empirically assessing the effect of increasing M on the bias, we find diminishing returns using three standard GEC measures (M^2 , accuracy and GLEU), underscoring the difficulty in this approach.

Similar trends are found when conducting such experiments to Text Simplification (TS) (§4). Specifically we show that (1) the distribution of valid simplifications for a given sentence is long-tailed; (2) common measures for TS dramatically under-estimate performance; (3) additional references alleviate this under-prediction.

To recap, we find that the LCB hinders the reliability of RBMs for GEC, and incentivizes systems developed to optimize these measures not to correct. LCB cannot be overcome by re-scaling or increasing M in any feasible range.

2 Coverage in RBMs

We begin by formulating a methodology for studying the distribution of valid corrections for a sentence (§2.1), and then turn to assessing the effect inadequate coverage has on common RBMs (§2.2). Finally, we compare human and system scores by common RBMs (§2.3).

Notation. We assume each ungrammatical sentence x has a set of valid corrections $Correct_x$, and a discrete distribution \mathcal{D}_x over them, where $P_{\mathcal{D}_x}(y)$ for $y \in Correct_x$ is the probability a human annotator would correct x as y .

Let $X = x_1 \dots x_N$ be the evaluated set of source sentences and denote $\mathcal{D}_i := \mathcal{D}_{x_i}$. Each x_i is independently sampled from some distribution \mathcal{L} over input sentences, and is paired with M corrections $Y_i = \{y_i^1, \dots, y_i^M\}$, which are independently sampled from \mathcal{D}_i . Our analysis assumes a fixed number of references across sentences, but generalizing to sentence-dependent M is straightforward. The *coverage* of a reference set Y_i of size M for a sentence x_i is defined as $P_{y \sim \mathcal{D}_i}(y \in Y_i)$.

A system C is a function from input sentences to proposed corrections (strings). An evaluation measure is a function $f: X \times Y \times C \rightarrow \mathbb{R}$. We use the term “true measure” to refer to a measure’s output where the reference set includes all valid corrections, i.e., $\forall i: Y_i = Correct_i$.

Experimental Setup. We conduct all experiments on the NUCLE test dataset (Dahlmeier

et al., 2013). NUCLE is a parallel corpus of essays written by language learners and their corrected versions, containing 1414 essays and 50 test essays, each of about 500 words.

We evaluate all participating systems in the CoNLL 2014 shared task, in addition to three of the best performing systems on this dataset, a hybrid system (Rozovskaya and Roth, 2016), a phrase-based MT system (Junczys-Dowmunt and Grundkiewicz, 2016) and a neural network system (Xie et al., 2016). Appendix A lists system names and abbreviations.

2.1 Estimating the Corrections Distribution

Data. We turn to estimating the number of corrections per sentence, and their histogram. The experiments in the following section are run on a random sample of 52 short sentences from the NUCLE test data, i.e. with 15 words or less. Through the length restriction, we avoid introducing too many independent errors that may drastically increase the number of annotation variants (as every combination of corrections for these errors is possible), thus resulting in unreliable estimation for \mathcal{D}_x .

Proven effective in GEC and related tasks such as MT (Zaidan and Callison-Burch, 2011; Madnani et al., 2011; Post et al., 2012), we use crowdsourcing to sample from \mathcal{D}_x (see Appendix B). Aiming to judge grammaticality rather than fluency, we instructed the workers to correct only when necessary, not for styling. We begin by estimating the histogram of \mathcal{D}_x for each sentence, using the crowdsourced corrections. We use UNSEENEST (Zou et al., 2016), a non-parametric algorithm to estimate a discrete distribution in which the individual values do not matter, only their probability. UNSEENEST aims to minimize the “earthmover distance”, between the estimated histogram and the histogram of the distribution. Intuitively, if histograms are piles of dirt, UNSEENEST minimizes the amount of dirt moved times the distance it moved. UNSEENEST was originally developed and tested for estimating the histogram of variants a gene may have, including undiscovered ones, a setting similar to ours. Our manual tests of UNSEENEST with small artificially created datasets showed satisfactory results.¹

¹An implementation of UNSEENEST, the data we collected, the estimated distributions and efficient implementations of computations with Poisson binomial distributions can be found in <https://github.com/borgr/IBGEC>.

Our estimates show that most input sentences have a large number of infrequent corrections that account for much of the probability mass and a rather small number of frequent corrections. Table 1 presents the mean number of different corrections with frequency at least γ (for different γ s), and their total probability mass. For instance, 74.34 corrections account for 75% of the probability mass, each occurring with frequency $\geq 0.1\%$.

	Frequency Threshold (γ)			
	0	0.001	0.01	0.1
Variants	1351.24	74.34	8.72	1.35
Mass	1	0.75	0.58	0.37

Table 1: Estimating the distribution of corrections \mathcal{D}_x . The table presents the mean number of corrections per sentence with probability more than γ (top row), as well as their total probability mass (bottom row).

The high number of rare corrections raises the question of whether these can be regarded as noise. To test this we conducted another crowdsourcing experiment, where 3 annotators were asked to judge whether a correction produced in the first experiment, is indeed valid. We plot the validity of corrections against their frequencies, finding that frequency has little effect, where even the rarest corrections are judged valid 78% of the time. Details in Appendix C.

2.2 Under-estimation as a Function of M

After estimating the histogram of valid corrections for a sentence, we turn to estimating the resulting bias (LCB), for different M values. We study sentence-level accuracy, F -Score and GLEU.

Sentence-level Accuracy. Sentence-level accuracy is the percentage of corrections that exactly match one of the references. Accuracy is a basic, interpretable measure, used in GEC by, e.g., Rozovskaya and Roth (2010). It is also closely related to the 0-1 loss function commonly used for training in GEC (Chodorow et al., 2012; Rozovskaya and Roth, 2013).

Formally, given test sentences $X = \{x_1, \dots, x_N\}$, their references Y_1, \dots, Y_N and a system C , we define C 's accuracy to be

$$Acc(C; X, Y) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{C(x_i) \in Y_i}. \quad (1)$$

Note that C 's accuracy is, in fact, an estimate of C 's *true accuracy*, the probability to produce a valid correction for a sentence. Formally:

$$TrueAcc(C) = P_{x \sim L}(C(x) \in Correct_x). \quad (2)$$

The bias of $Acc(C; X, Y)$ for a sample of N sentences, each paired with M references is then

$$TrueAcc(C) - \mathbb{E}_{X, Y}[Acc(C; X, Y)] = \quad (3)$$

$$TrueAcc(C) - P(C(x) \in Y) = \quad (4)$$

$$P(C(x) \in Correct_x) \cdot \quad (5)$$

$$(1 - P(C(x) \in Y | C(x) \in Correct_x)) \quad (6)$$

We observe that the bias, denoted b_M , is not affected by N , only by M . As M grows, Y better approximates $Correct_x$, and b_M tends to 0.

In order to abstract away from the idiosyncrasies of specific systems, we consider an idealized learner, which, when correct, produces a valid correction with the same distribution as a human annotator (i.e., according to \mathcal{D}_x). Formally, we assume that, if $C(x) \in Correct_x$ then $C(x) \sim \mathcal{D}_x$. Hence the bias b_M (Eq. 6) can be re-written as

$$P(C(x) \in Correct_x) \cdot (1 - P_{Y \sim \mathcal{D}_x^M, y \sim \mathcal{D}_x}(y \in Y)).$$

We will henceforth assume that C is perfect (i.e., its *true accuracy* is 1). Note that assuming any other value for C 's *true accuracy* would simply scale b_M by that accuracy. Similarly, assuming only a fraction p of the sentences require correction scales b_M by p .

We estimate b_M empirically using its empirical mean on our experimental corpus:

$$\hat{b}_M = 1 - \frac{1}{N} \sum_{i=1}^N P_{Y \sim \mathcal{D}_i^M, y \sim \mathcal{D}_i}(y \in Y).$$

Using the UNSEENEST estimations of \mathcal{D}_i , we can compute \hat{b}_M for any size of Y_i (M). However, as this is highly computationally demanding, we estimate it using sampling. Specifically, for every $M = 1, \dots, 20$ and x_i , we sample Y_i 1000 times (with replacement), and estimate $P(y \in Y_i)$ as the covered probability mass $P_{\mathcal{D}_i}\{y : y \in Y_i\}$. Based on that we compute the accuracy distribution and expectation (see Appendix D).

We repeated all our experiments where Y_i is sampled without replacement, and find similar trends with a faster increase in accuracy reaching over 0.47 with $M = 10$.

Figure 1a presents the expected accuracy of a perfect system (i.e., $1 - \hat{b}_M$) for different M s. Results show that even for M values which are much larger than the standard (e.g., $M = 20$), expected




Figure 1: The score obtained by perfect systems according to GEC accuracy (1a), GEC F-score and GLEU (1b). Figure 1c reports TS experimental results, namely the score of a perfect and lucky perfect system using SARI, and a perfect system using MAX-SARI. The y-axis corresponds to the measure values, and the x-axis to the number of references M . For bootstrapping experiments points are paired with a confidence interval ($p = .95$).

accuracy is only around 0.5. As M increases, the contribution of each additional correction diminishes sharply (the slope is 0.004 for $M = 20$).

We also experiment with a more relaxed measure, *Exact Index Match*, which is only sensitive to the identity of the changed words and not to what they were changed to. Formally, two corrections c and c' over a source sentence x match if for their word alignments with the source (computed as above) $a : \{1, \dots, |x|\} \rightarrow \{1, \dots, |c|, Null\}$ and $a' : \{1, \dots, |x|\} \rightarrow \{1, \dots, |c'|, Null\}$, it holds that $c_{a(i)} \neq x_i \Leftrightarrow c'_{a'(i)} \neq x_i$, where $c_{Null} = c'_{Null}$. Results, while somewhat higher, are still only 0.54 with $M = 10$. (Figure 1a)

F-Score. While accuracy is commonly used as a loss function for training GEC systems, F_α -score is standard for evaluating system performance. The score is computed in terms of *edit* overlap between edits that constitute a correction and ones that constitute a reference, where edits are substring replacements to the source. We use the standard M^2 scorer (Dahlmeier and Ng, 2012), which defines edits optimistically, maximizing over all possible annotations that generate the correction from the source. Since our crowdsourced corrections are not annotated for edits, we produce edits to the reference heuristically.

The complexity of the measure prohibits an analytic approach (Yeh, 2000). We instead use bootstrapping to estimate the bias incurred by not being able to exhaustively enumerate the set of valid corrections. As with accuracy, in order to avoid confounding our results with system-specific biases, we assume the evaluated system is perfect and sample its corrections from the human distribution of corrections \mathcal{D}_x .

Concretely, given a value for M and for N , we uniformly sample from our experimental corpus source sentences x_1, \dots, x_N , and M corrections for each Y_1, \dots, Y_N (with replacement). Setting a realistic value for N in our experiments is important for obtaining comparable results to those obtained on the NUCLE corpus (see §2.3), as the expected value of F -score depends on N and the number of sentences that do not need correction (N_{cor}). Following the statistics of NUCLE’s test set, we set $N = 1312$ and $N_{cor} = 136$.

Bootstrapping is carried out by the accelerated bootstrap procedure (Efron, 1987), with 1000 iterations. We also report confidence intervals ($p = .95$), computed using the same procedure.

Results (Figure 1b) again show the insufficiency of commonly-used M values for reliably estimating system performance. For instance, the $F_{0.5}$ -score for our perfect system is only 0.42 with $M = 2$. The saturation effect, observed for accuracy, is even more pronounced in this setting.

GLEU. We repeat the procedure using the mean GLEU sentence score (Figure 1b), which was shown to better correlate with human judgments than M^2 (Napolet et al., 2016). Results are about 2% higher than M^2 ’s with a similar saturation effect. Sakaguchi et al. (2016) observed a similar effect when evaluating against fluency-oriented references; this has led them to assume that saturation is due to covering most of the probability mass, which we now show is not the case.²

² We do not experiment with I-measure (Felice and Briscoe, 2015), as its run-time is prohibitively high for experimenting with bootstrapping that requires many applications of the measure (Choshen and Abend, 2018a), and as empirical validation studies showed that it has a low correlation with human judgments (Sakaguchi et al., 2016).

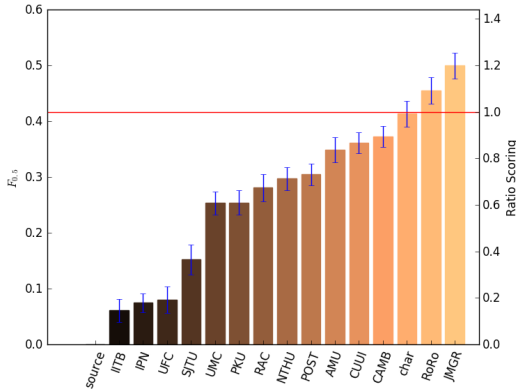


Figure 2: $F_{0.5}$ values with $M = 2$ for different systems, including confidence interval ($p = .95$). The left-most column (“source”) presents the F -score of a system that doesn’t make any changes to the source sentences. In red is human performance. See §2 for a legend of the systems.

2.3 Human and System Performance

The bootstrapping method for computing the significance of the F -score (§2.2) can also be used for assessing the significance of the differences in system performance reported in the literature. We compute confidence intervals of different systems on the NUCLE test data ($M = 2$).

Results (Figure 2) present mixed trends: some differences between previously reported F -scores are indeed significant and some are not. For example, the best performing system is significantly better than all but the second one.

Considering the F -score of the best-performing systems, and comparing them to the F -score of a perfect system with $M = 2$ (in accordance with systems’ reported results), we find that their scores are comparable, where the systems RoRo and JMGR surpass a perfect system’s F -score. Similar experiments with GLEU show that the two systems obtain comparable or superior performance to humans on this measure as well.

2.4 Discussion

In this section we have established that (1) as systems can surpass human performance on RBMs, re-scaling cannot be used to overcome the LCB, and that (2) as the distribution of valid corrections is long-tailed, the number of references needed for reliable RBMs is exceedingly high. Indeed, an average sentence has hundreds or more valid low-probability corrections, whose total probability mass is substantial. Our analysis with Exact Index Match suggests that similar effects are applicable to Grammatical Error Detection as well. The

proposal of Sakaguchi et al. (2016), to emphasize fluency over grammaticality in reference corrections, only compounds this problem, as it results in a larger number of valid corrections.

3 Implications of the LCB

We discuss the adverse effects of LCB not only on the reliability of RBMs, but on the development of GEC systems. We argue that evaluation with inadequate reference coverage incentivizes systems to under-correct, and to mostly target errors that have few valid corrections (closed-class). We first show that low coverage can lead to under-correction (§3.1), then show that modern systems make far fewer corrections to the source, compared to humans (§3.2). §3.3 shows that increasing the number of references can alleviate this effect. §3.4 shows that open-class errors are more likely to be under-corrected than closed-class ones.

3.1 Motivating Analysis

For simplicity, we abstract away from the details of the learning model and assume that systems attempt to maximize an objective function, over some training or development data. We assume maximization is achieved by iterating over the samples, as with the Perceptron or SGD.

Assume the system is faced with a phrase it predicts to be ungrammatical. Assume p_{detect} is the probability this prediction is correct, and $p_{correct}$ is the probability it is able to predict a valid correction for this phrase (including correctly identifying it as erroneous). Finally, assume evaluation is against M references with coverage $p_{coverage}$ (the probability that a valid correction will be found among M randomly sampled references).

We will now assume that the system may either choose to correct with the correction it finds the most likely or not at all. If it chooses not to correct, its probability of being rewarded (i.e., its output is in the reference set) is $(1 - p_{detect})$. Otherwise, its probability of being rewarded is $p_{correct} \cdot p_{coverage}$. A system is disincentivized from altering the phrase in cases where:

$$p_{correct} \cdot p_{coverage} < 1 - p_{detect} \quad (7)$$

We expect Condition (7) to frequently hold in cases that require non-trivial changes, which are characterized both by low $p_{coverage}$ (as non-trivial changes are often open-class), and by lower system performance.

Corrector	Sentence
Source	This is especially to people who are overseas.
CHAR, UMC, JMGR	This is especially for people who are overseas.
IPN	This is especially to peoples who are overseas.
CUUI	This is especially to the people who are overseas.
NUCLE _A	This is especially true for people who are overseas.
NUCLE _B	This is especially relevant to people who are overseas.

Table 2: Example for a sentence and proposed corrections by different systems (top part) and by the two NUCLE annotators (bottom part). Systems not mentioned in the table retain the source. No system produces a new word as needed. The two references differ in their corrections.

Precision-oriented measures (e.g., $F_{0.5}$) penalize invalidly correcting more harshly than not correcting an ungrammatical sentence. In these cases, Condition (7) should be written as

$$p_{correct} \cdot p_{coverage} - (1 - p_{correct} \cdot p_{coverage}) \alpha < 1 - p_{detect}$$

where α is the ratio between the penalty for introducing a wrong correction and the reward for a valid correction. The condition is even more likely to hold with such measures.

3.2 Under-correction in GEC Systems

In this section we compare the prevalence of changes made to the source by the systems, to their prevalence in the NUCLE references. To strengthen our claim, we exclude all non-alphanumeric characters, both within tokens or as separate tokens. See Table 2 for an example.

We consider three types of divergences between the source and the reference. First, we measure the extent to which *words* were changed: altered, deleted or added. To do so, we compute word alignment between the source and the reference, casting it as a weighted bipartite matching problem. Edge weights are assigned to be the token edit distances.³ Following word alignment, we define WORDCHANGE as the number of aligned words and unaligned words changed. Second, we quantify word *order* differences using Spearman’s ρ between the order of the words in the source sentence and the order of their corresponding-aligned words in the correction. $\rho = 0$ where the word

³Aligning words in GEC is much simpler than in MT, as most of the words are unchanged, deleted fully, added, or changed slightly.

order is uncorrelated, and $\rho = 1$ where the orders exactly match. We report the average ρ over all source sentence pairs. Third, we report how many source sentences were split and how many concatenated by the reference and by the systems. One annotator was arbitrarily selected for the figures.

Results. Results (Figure 3) show that humans make considerably more changes than systems according to all measures of under-correction, both in terms of the number of sentences modified and the number of modifications within them. Differences are often an order of magnitude large. For example, 36 reference sentences include 6 word changes, where the maximal number of sentences with 6 word changes by any system is 5. We find similar trends on the references of the TreeBank of Learner English (Yannakoudakis et al., 2011).

3.3 Higher M Alleviates Under-correction

This section reports an experiment for determining whether increasing the number of references in training indeed reduces under-correction. There is no corpus available with multiple references which is large enough for re-training a system. Instead, we simulate such a setting with an oracle reranking approach, and test whether the availability of increasingly more training references reduces a system’s under-correction.

Concretely, given a set of sentences, each paired with M references, a measure and a system’s k -best list, we define an oracle re-ranker that selects for each sentence the highest scoring correction. As a test case, we use the RoRo system with $k = 100$, and apply it to the largest available language learner corpus which is paired with a substantial amount of GEC references, namely the NUCLE test corpus. We use the standard F -score as the evaluation measure, examining the under-correction of the oracle re-ranker for different M values, averaging over the 1312 samples of M references from the available set of ten references provided by Bryant and Ng (2015).

As the argument is not trivial, we turn to explaining why decreased under-correction with an increase in M indicates that tuning against a small set of references (low coverage) yields under-correction. Assume an input sentence with some sub-string e . There are three cases: (1) e is an error, (2) e is valid but there are valid references that alter it, (3) e is uniquely valid. In case (3) or-




Figure 3: The prevalence of changes in system outputs and in the NUCLE reference. The top figure presents the number of sentences (heat) for each amount of word changes (x-axis; measured by WORDCHANGE) done by the outputs and the reference (y-axis). The middle figure presents the percentage of sentence pairs (y-axis) where the Spearman ρ values do not exceed a certain threshold (x-axis). The bottom figure presents the counts of source sentences (y-axis) concatenated (right bars) or split (left bars) by the references (striped column) and the outputs (coloured columns). See Appendix A for a legend of the systems. Under all measures, the gold standard references make substantially more changes to the source sentences than any of the systems, in some cases an order of magnitude more.

	L_{val} empty	L_{val} not empty	
		e valid	e error
Small M	0	$P_Y(\bar{e}, L_{val})$	$P_Y(L_{val})$
Large M	0	0	1
Correction Rate	=	↓	↑

Table 3: The expected effect of oracle re-ranking on under-correction. Values represent the probability of altering a substring of the input e , which is a proxy to the expected correction rate. L_{val} is the valid alterations in the k -best list. $P_Y(L_{val})$ is the probability that a valid correction from the list is also in the reference set Y , $P_Y(\bar{e}, L_{val})$ is the probability that, in addition, the reference that keeps e is not in Y . When M increases, the expected correction rate is expected to increase only if e is an error and a valid correction of it is found in the k -best list.

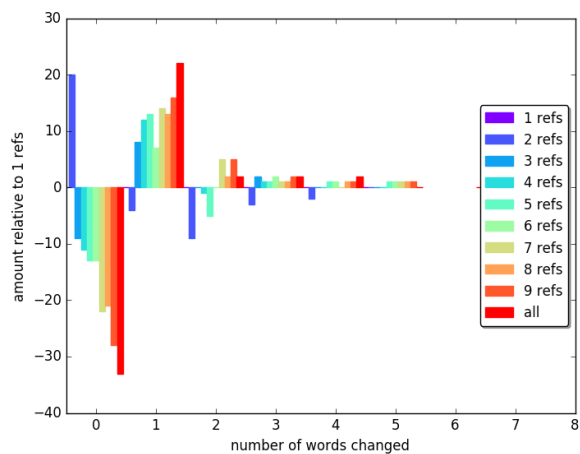


Figure 4: The amount of sentences (y-axis) with a given number of words changed (x-axis) following oracle reranking with different M values (column colors), where the amount for $M = 1$ is subtracted from them. All references are randomly sampled except the “all” column that contains all ten references. In conclusion, tuning against additional references indeed reduces under-correction.

acle re-ranking has no effect and can be ignored. The corrections in the k -best list can then be partitioned to those that keep e as it is; those that invalidly alter e ; and those that validly alter e .

Table 3 presents the probability that e will be altered in the different cases. Analysis shows that under-correction is likely to decrease with M only in the case where e is an error and the k -best list contains a valid correction of it. Whenever the reference allows both keeping e and altering e , the re-ranker selects keeping e .

Indeed, our experimental results show that word changes increase with M (Figure 4), indicating that low coverage may play a role in the observed tendency of GEC systems to under-correct. No significant difference is found for word order.

3.4 Under-correction by Error Types

In this section we study the prevalence of under-correction according to edit types, finding that open-class types of errors (such as replacing a word with another word) are more starkly under-corrected, than closed-class errors. Evaluating with low coverage RBMs does not incentivize systems to address open-class errors (in fact, it disincentivizes them to). Therefore, even if LCB is not the cause for this trend, current evaluation procedures may perpetuate it.

We use the data of Bryant et al. (2017), which automatically assigned types to each edit in the output of all CoNLL 2014 systems on the NUCLE test set. As a measure of under-correction tendency, we take the ratio between the mean number of corrections produced by the systems and by the references. We note that this analysis does not consider whether the predicted correction is valid or not, but only how many of the errors of each type the systems attempted to correct.

We find that all edit types are under-predicted on average, but that the least under-predicted ones are mostly closed-class types. Concretely, the top quarter of error types consists of orthographical errors, plurality inflection of nouns, adjective inflections to superlative or comparative forms and determiner selection. The bottom quarter includes the categories verb selection, noun selection, particle/preposition selection, pronoun selection, and the type OTHER, which is a residual category. The only exception to this regularity is the closed-class punctuation selection type, which is found in the lower quarter. See Appendix E.

This trend cannot be explained by assuming that common error types are targeted more. Indeed, error type frequency is slightly negatively correlated with the under-correction ratio ($\rho=-0.29$ p-value=0.16). A more probable account of this effect is the disincentive of GEC systems to correct open-class error types, for which even valid corrections are unlikely to be rewarded.

4 Similar Effects on Simplification

We now turn to replicating our experiments on Text Simplification (TS). From a formal point of view, evaluation of the tasks is similar: the output is obtained by making zero or more edits to the source. RBMs are the standard for TS evaluation, much like they are in GEC.

Our experiments on TS demonstrate that simi-

lar trends recur in this setting as well. The tendency of TS systems to under-predict changes to the source has already been observed by previous work (Alva-Manchego et al., 2017), showing that TS systems under-predict word additions, deletions, substitutions, and sequence shifts (Zhang and Lapata, 2017), and have low edit distance from the source (Narayan and Gardent, 2016). Our experiments show that LCB may account for this under-prediction. Concretely, we show that (1) the distribution of valid references for a given sentence is long-tailed; (2) common evaluation measures suffer from LCB, taking SARI (Xu et al., 2016) as an example RBM (similar trends are obtained with Accuracy); (3) under-prediction is alleviated with M in oracle re-ranking experiments.

We crowd-sourced 2500 reference simplifications for 47 sentences, using the corpus and the annotation protocol of Xu et al. (2016), and applying UNSEENEST to estimate \mathcal{D}_x (Appendix B). Table 4 shows that the expected number of references is even greater in this setting.

Assessing the effect of M on SARI, we find that SARI diverges from Accuracy and F -score in that its multi-reference version is not a maximum over the single-reference scores, but some combination of them. This can potentially increase coverage, but it also leads to an unintuitive situation: an output identical to a reference does not receive a perfect score, but rather the score depends on how similar the output is to the other references. A more in-depth analysis of SARI’s handling of multiple references is found in Appendix F. In order to neutralize this effect of SARI, we also report results with MAX-SARI, which coincides with SARI on $M = 1$, and is defined as the maximum single-reference SARI score for $M > 1$.

Figure 1c presents the coverage of SARI and MAX-SARI of a perfect TS system that selects a random correction from the estimated distribution of corrections using the same bootstrapping protocol as in §2.1. We also include the SARI score of a “lucky perfect” system, that randomly selects one of the given references (the MAX-SARI score for such a system is 1). Results show that SARI has a coverage of about 0.45, and that this score is largely independent of M . The score of predicting one of the available references drops with the number of references, indicating that SARI scores may not be comparable across different M values.

We therefore restrict oracle re-ranking experi-

	Frequency Threshold (γ)			
	0	0.001	0.01	0.1
Variants	2636.29	111.19	4.68	0.13
Mass	1	0.42	0.14	0.02

Table 4: Estimating the distribution of simplifications \mathcal{D}_x . The table presents the mean number of simplifications per sentence with probability more than γ (top row), as well as their total probability mass (bottom row).

ments to MAX-SARI, conducting re-ranking experiments on k -best lists in two settings: Moses (Koehn et al., 2007) with $k = 100$, and a neural model (Nisioi et al., 2017) with $k = 12$. Our results indeed show that under-prediction is alleviated with M in both settings. For example, the least under-predicting model (the neural one) did not change 50 sentences with $M = 1$, but only 29 weren’t changed with $M = 8$. See Appendix G.

5 Conclusion

We argue that using low-coverage reference sets has adverse effects on the reliability of reference-based evaluation, with GEC and TS as a test case, and consequently on the incentives offered to systems. We further argue that these effects cannot be overcome by re-scaling or increasing the number of references in a feasible way. The paper makes two methodological contributions to the monolingual translation evaluation literature: (1) a methodology for evaluating evaluation measures by the scores they assign a perfect system, using a bootstrapping procedure; (2) a methodology for assessing the distribution of valid monolingual translations. Our findings demonstrate how these tools can help characterize the biases of existing systems and evaluation measures. We believe our findings and methodologies can be useful for similar tasks such as style conversion and automatic post-editing of raw MT outputs.

We note that the LCB further jeopardizes the reliability of common validation experiments for RBMs, that assess the correlation between human and measure rankings of system outputs (Grundkiewicz et al., 2015). Indeed, if outputs all similarly under-correct, correlation studies will not be affected by whether an RBM is sensitive to under-correction. Therefore, the tendency of RBMs to reward under-correction cannot be detected by such correlation experiments (cf. Choshen and Abend, 2018a).

Our results underscore the importance of de-

veloping alternative evaluation measures that transcend n -gram overlap, and use deeper analysis tools, e.g., by comparing the semantics of the reference and the source to the output (cf. Lo and Wu, 2011). Napoles et al. (2016) have made progress towards this goal in proposing a reference-less grammaticality measure, using Grammatical Error Detection tools, as did Asano et al. (2017), who added a fluency measure to the grammaticality. In a recent project (Choshen and Abend, 2018b), we proposed a complementary measure that measures the semantic faithfulness of the output to the source, in order to form a combined semantic measure that bypasses the pitfalls of low coverage.

Acknowledgments

This work was supported by the Israel Science Foundation (grant No. 929/17), and by the HUJI Cyber Security Research Center in conjunction with the Israel National Cyber Bureau in the Prime Minister’s Office. We thank Nathan Schneider, Courtney Napoles and Joel Tetreault for helpful feedback.

References

- Fernando Alva-Manchego, Joachim Bingel, Gustavo Paetzold, Carolina Scarton, and Lucia Specia. 2017. [Learning how to simplify from explicit labeling of complex-simplified text pairs](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 295–305, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Hiroki Asano, Tomoya Mizumoto, and Kentaro Inui. 2017. Reference-based metrics can be replaced with reference-less metrics in evaluating grammatical error correction systems. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, volume 2, pages 343–348.
- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. [Automatic annotation and evaluation of error types for grammatical error correction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.
- Christopher Bryant and Hwee Tou Ng. 2015. How far are we from fully automatic high quality grammatical error correction? In *ACL (1)*, pages 697–707.
- Martin Chodorow, Markus Dickinson, Ross Israel, and Joel R Tetreault. 2012. Problems in evaluating grammatical error detection systems. In *COLING*, pages 611–628. Citeseer.

- Leshem Choshen and Omri Abend. 2018a. Automatic metric validation for grammatical error correction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Leshem Choshen and Omri Abend. 2018b. Referenceless measure of faithfulness for grammatical error correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Daniel Dahlmeier and Hwee Tou Ng. 2012. Better evaluation for grammatical error correction. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572. Association for Computational Linguistics.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. Building a large annotated corpus of learner english: The nus corpus of learner english. In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 22–31.
- Markus Dreyer and Daniel Marcu. 2012. Hyter: Meaning-equivalent semantics for translation evaluation. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 162–171. Association for Computational Linguistics.
- Bradley Efron. 1987. Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397):171–185.
- Mariano Felice and Ted Briscoe. 2015. Towards a standard evaluation method for grammatical error detection and correction. In *HLT-NAACL*, pages 578–587.
- Roman Grundkiewicz, Marcin Junczys-Dowmunt, Edward Gillian, et al. 2015. Human evaluation of grammatical error correction systems. In *EMNLP*, pages 461–470.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. Phrase-based machine translation is state-of-the-art for automatic grammatical error correction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1546–1556, Austin, Texas. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180.
- Chi-kiu Lo and Dekai Wu. 2011. Meant: an inexpensive, high-accuracy, semi-automatic metric for evaluating translation utility via semantic frames. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 220–229. Association for Computational Linguistics.
- Nitin Madnani, Joel Tetreault, Martin Chodorow, and Alla Rozovskaya. 2011. They can help: Using crowdsourcing to improve the evaluation of grammatical error detection systems. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, pages 508–513. Association for Computational Linguistics.
- Inderjeet Mani. 2009. Summarization evaluation: an overview. In *Proceedings of the NTCIR Workshop*, volume 2.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. Ground truth for grammatical error correction metrics. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, volume 2, pages 588–593.
- Courtney Napoles, Keisuke Sakaguchi, and Joel Tetreault. 2016. There’s no comparison: Referenceless evaluation metrics in grammatical error correction. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2109–2115. Association for Computational Linguistics.
- Shashi Narayan and Claire Gardent. 2016. [Unsupervised sentence simplification using deep semantics](#). In *Proceedings of the 9th International Natural Language Generation conference*, pages 111–120, Edinburgh, UK. Association for Computational Linguistics.
- Sergiu Nisioi, Sanja Štajner, Simone Paolo Ponzetto, and Liviu P Dinu. 2017. Exploring neural text simplification models. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 85–91.
- Matt Post, Chris Callison-Burch, and Miles Osborne. 2012. Constructing parallel corpora for six indian languages via crowdsourcing. In *Proceedings of the Seventh Workshop on Statistical Machine Translation*, pages 401–409. Association for Computational Linguistics.
- Alla Rozovskaya and Dan Roth. 2010. Annotating esl errors: Challenges and rewards. In *Proceedings of the NAACL HLT 2010 fifth workshop on innovative use of NLP for building educational applications*,

- pages 28–36. Association for Computational Linguistics.
- Alla Rozovskaya and Dan Roth. 2013. Joint learning and inference for grammatical error correction. *Urbana*, 51:61801.
- Alla Rozovskaya and Dan Roth. 2016. Grammatical error correction: Machine translation and classifiers. In *Proc. of ACL*, pages 2205–2215.
- Keisuke Sakaguchi, Courtney Napoles, Matt Post, and Joel Tetreault. 2016. Reassessing the goals of grammatical error correction: Fluency instead of grammaticality. *Transactions of the Association for Computational Linguistics*, 4:169–182.
- Joel R Tetreault and Martin Chodorow. 2008. Native judgments of non-native usage: Experiments in preposition error detection. In *Proceedings of the Workshop on Human Judgements in Computational Linguistics*, pages 24–32. Association for Computational Linguistics.
- Ziang Xie, Anand Avati, Naveen Arivazhagan, Dan Jurafsky, and Andrew Y Ng. 2016. Neural language correction with character-based attention. *arXiv preprint arXiv:1603.09727*.
- Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3:283–297.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. [Optimizing statistical machine translation for text simplification](#). *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading esol texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 180–189. Association for Computational Linguistics.
- Alexander Yeh. 2000. More accurate tests for the statistical significance of result differences. In *Proceedings of the 18th conference on Computational linguistics-Volume 2*, pages 947–953. Association for Computational Linguistics.
- Omar F Zaidan and Chris Callison-Burch. 2011. Crowdsourcing translation: Professional quality from non-professionals. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1220–1229. Association for Computational Linguistics.
- Xingxing Zhang and Mirella Lapata. 2017. [Sentence simplification with deep reinforcement learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 584–594, Copenhagen, Denmark. Association for Computational Linguistics.
- James Zou, Gregory Valiant, Paul Valiant, Konrad Karczewski, Siu On Chan, Kaitlin Samocha, Monkol Lek, Shamil Sunyaev, Mark Daly, and Daniel G MacArthur. 2016. Quantifying unobserved protein-coding variants in human populations provides a roadmap for large-scale sequencing projects. *Nature Communications*, 7.