

On the Limitations of Reference-Free Evaluations of Generated Text

Daniel Deutsch,^{*†} Rotem Dror,[‡] and Dan Roth[‡]

[†]Google Research

[‡]University of Pennsylvania

dandeutsch@google.com

{rtmdrr, danroth}@seas.upenn.edu

Abstract

There is significant interest in developing evaluation metrics which accurately estimate the quality of generated text without the aid of a human-written reference text, which can be time consuming and expensive to collect or entirely unavailable in online applications. However, in this work, we demonstrate that these reference-free metrics are inherently biased and limited in their ability to evaluate generated text, and we argue that they should not be used to measure progress on tasks like machine translation or summarization. We show how reference-free metrics are equivalent to using one generation model to evaluate another, which has several limitations: (1) the metrics can be optimized at test time to find the approximate best-possible output, (2) they are inherently biased toward models which are more similar to their own, and (3) they can be biased against higher-quality outputs, including those written by humans. Therefore, we recommend that reference-free metrics should be used as diagnostic tools for analyzing and understanding model behavior instead of measures of how well models perform a task, in which the goal is to achieve as high of a score as possible.¹

1 Introduction

Automatically evaluating the quality of generated texts is essential for the development of natural language generation systems. The most common type of evaluation for generation tasks such as machine translation (MT) and summarization is done with reference-based automatic metrics, which evaluate a text by comparing it to a gold-standard reference text, usually written by humans (Papineni et al., 2002; Lin, 2004; Zhang et al., 2020; Sellam et al.,

		Prism-src (↑)
Source	Doch er ist nicht krank, er hat nur einen mächtigen Kater .	
Reference	But he is not ill, he only has quite a hangover .	-1.6
Candidate	But he is not sick, he has only one powerful cat .	-0.4
Source	Und mit Mann und Maus gegen Mainz verteidigt.	
Reference	And threw everything they had into our defense.	-4.8
Candidate	And defended with man and mouse against Mainz.	-0.4

Figure 1: Here, Prism-src was optimized to generate the candidate translations. They are clearly wrong (*Kater* means both “cat” and “hangover”; *mit Mann und Maus* is an expression that means “with all means available”), but have better Prism-src scores than the references. Comparing systems with reference-free metrics will favor systems that are more similar to the metrics’ underlying models rather than higher quality output.

2020; Deutsch et al., 2021a; Zhang and Bansal, 2021, *inter alia*).

Reference texts can be expensive to collect or are entirely unavailable when there is a need to estimate the quality of text in real time, so there is an increased interest in developing automatic metrics that do not use references to evaluate text, commonly referred to as reference-free metrics (Louis and Nenkova, 2013; Fonseca et al., 2019; Scialom et al., 2019, 2021; Vasilyev et al., 2020; Rei et al., 2021, *inter alia*). While these metrics do not always achieve performance parity with their reference-based counterparts, their high correlations to human judgments suggest that reference-free evaluation is a promising direction of future research (Fonseca et al., 2019).²

However, in this work, we demonstrate that reference-free evaluation metrics have inherent limitations and argue that they should not be used

¹https://cogcomp.seas.upenn.edu/page/publication_view/991

^{*}Work done while at the University of Pennsylvania

²Some reference-free metrics actually already out-perform reference-based metrics (Freitag et al., 2021).

to measure progress on tasks, even in domains in which no reference texts are available. Central to our argument is the idea that because reference-free metrics evaluate text using the same input provided to the generation models, they are either explicitly or implicitly using an underlying generation model to evaluate other models (§2). There are several implications of this, which we explore through an analysis of three reference-free evaluation metrics, Prism-src (Thompson and Post, 2020) and COMET-QE (Rei et al., 2021) for MT and Quest-Eval (Scialom et al., 2021) for summarization.

First, the metrics’ underlying models will achieve the best possible metric score by definition. Therefore, the “perfect” model is already known, and we show that it is possible to define simple approximate inference algorithms which use these models to find the approximate best output according to the metrics (§4, §5.1).

Then, the metrics have inherent, undesirable biases that originate from their underlying models. Not only do they favor the underlying models’ outputs, but they are also biased toward outputs from models which are similar to their own, and biased against higher-quality outputs, such as those written by humans (Fig. 1, §5.2, §5.3). Thus, if they were used as primary evaluation methods for a task, they would encourage other models to be more similar to their own and less human-like, an undesirable property of an evaluation metric.

Our recommendation is that reference-free metrics should not be used as methods for measuring progress on generation tasks such as MT, in which the goal is to achieve the highest possible value of the metric. Instead, they are better suited to be diagnostic statistics for analyzing model behavior with the understanding that they are inherently limited and biased (§6).

The contributions of this work include: (1) insight on the equivalence of reference-free metrics and generation models, (2) a demonstration that reference-free metrics’ values can be optimized at test time to achieve high-scoring outputs, and (3) an analysis that reveals reference free metrics’ inherent biases and limitations.

2 Reference-Free Metrics as Models

Conditional text generation models can be viewed as a function $\theta(\cdot)$ which scores an output text $\mathbf{y} \in \mathcal{Y}$ for some input text \mathbf{x} . Then $\theta(\cdot)$ is used in conjunction with an inference procedure $f_\theta(\cdot)$ to

find the best output at test time.³

$$\theta(\mathbf{x}, \mathbf{y}) \rightarrow \mathbb{R} \quad (1)$$

$$f_\theta(\mathbf{x}) = \arg \max_{\mathbf{y} \in \mathcal{Y}} \theta(\mathbf{x}, \mathbf{y}) \quad (2)$$

For instance, $\theta(\cdot)$ could be a learned sequence-to-sequence model and $f_\theta(\cdot)$ could be beam search.

The output of $f_\theta(\cdot)$, denoted $\hat{\mathbf{y}}$, is typically evaluated by some automatic metric \mathcal{M} . Reference-based metrics do this by scoring $\hat{\mathbf{y}}$ using some gold-standard text \mathbf{y}^* (which is not available to the model during inference) and the input \mathbf{x} (which is not always used). For instance, $\mathcal{M}_{\text{Ref-Based}}$ could calculate a BLEU score (Papineni et al., 2002) between the output translation $\hat{\mathbf{y}}$ and the gold translation \mathbf{y}^* .

$$\mathcal{M}_{\text{Ref-Based}}(\mathbf{x}, \hat{\mathbf{y}}, \mathbf{y}^*) \rightarrow \mathbb{R} \quad (3)$$

In contrast, reference-free metrics calculate a score for $\hat{\mathbf{y}}$ without \mathbf{y}^* :

$$\mathcal{M}_{\text{Ref-Free}}(\mathbf{x}, \hat{\mathbf{y}}) \rightarrow \mathbb{R} \quad (4)$$

Such metrics include the three analyzed in this work, namely, Prism-src (Thompson and Post, 2020), COMET-QE (Rei et al., 2021), and Quest-Eval (Scialom et al., 2021).

Because $\theta(\cdot)$ and $\mathcal{M}_{\text{Ref-Free}}$ are both functions of only \mathbf{x} and \mathbf{y} (equivalently $\hat{\mathbf{y}}$), $\mathcal{M}_{\text{Ref-Free}}$ itself can be viewed as a conditional generation model. For some metrics, such as Prism-src, this is explicitly stated, whereas others are implicitly making this assumption. This is not the case for reference-based metrics since they additionally require \mathbf{y}^* as input.

Since reference-free metrics are equivalent to generation models, there must exist some inference procedure which finds the best output text under the metric, denoted $g_{\mathcal{M}_{\text{Ref-Free}}}(\cdot)$:

$$g_{\mathcal{M}_{\text{Ref-Free}}}(\mathbf{x}) = \arg \max_{\mathbf{y} \in \mathcal{Y}} \mathcal{M}_{\text{Ref-Free}}(\mathbf{x}, \mathbf{y}) \quad (5)$$

Computing $g_{\mathcal{M}_{\text{Ref-Free}}}(\cdot)$ may be computationally expensive because $\mathcal{M}_{\text{Ref-Free}}$ may not support efficient inference. However, the inference procedure does always exist, and will return the best possible output according to the reference-free metric by definition.

We explore the implications of using a model to evaluate other models by analyzing the behavior of three different reference-free evaluation metrics on two text generation tasks, MT and summarization.

³In practice, $f_\theta(\cdot)$ finds the *approximate* best output, not the global maximum of $\theta(\cdot)$.

3 Analysis Setup

Here, we discuss the datasets and metrics used in our analysis of reference-free metrics.

Datasets Our MT experiments are run on the data collected for the WMT’19 metrics shared task (Ma et al., 2019), which includes reference translations and human-judged model outputs for 10 to 20 translation systems across 18 language pairs.

The summarization experiments use the SummEval (Fabbri et al., 2021) and REALSumm (Bhandari et al., 2020) datasets, which consist of reference summaries and human-judged model outputs for 16 and 25 summarization models, respectively, collected from the CNN/DailyMail dataset (Nallapati et al., 2016).

Prism-src Prism-src is a reference-free evaluation translation metric that scores a translated text according to the log-probability of the translation conditioned on the original source text under a learned sequence-to-sequence translation model (Thompson and Post, 2020). The model is a multi-lingual MT model, meaning it was trained using many different language pairs, so the same learned parameters can be used to score translations in various languages.

COMET-QE COMET-QE (Rei et al., 2021) is a modification of the learned reference-based MT evaluation metric COMET (Rei et al., 2020). COMET embeds the candidate translation, source text, and reference translation using a cross-lingual encoder, creates a pooled featured representation using the three encodings, and trains the model end-to-end to predict human judgments of the quality of the candidate translation. COMET-QE uses the same architecture to predict a score for the candidate translation but only uses the candidate translation and source text to create the pooled feature representation, and is therefore reference-free.

QuestEval Scialom et al. (2021) proposed a reference-free summarization metric called QuestEval which generates QA pairs from both the source document and generated summary then scores the summary based on the proportion of those pairs which are answered correctly in the opposite text. The metric optionally includes a step in which the QA pairs generated from the source document are weighted based on a learned query weighting model. The query weighter was trained to predict the probability that a question is answered in

the CNN/DailyMail reference summaries using a pre-trained QA model. We use the query weighter in our experiments since it improved the performance of QuestEval in Scialom et al. (2021).

Reference-Based Metrics We analyze the reference-free metrics with respect to various reference-based metrics which have been demonstrated to have strong correlations to human judgments of translation/summary quality. BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) compare the two texts using n -gram overlap statistics. BERTScore calculates a quality score based on how similar the reference and candidate texts’ BERT (Devlin et al., 2019) embeddings are (Zhang et al., 2020). QAEval is a QA-based metric for summarization, which generates wh-questions from the reference summary and calculates a score for the candidate summary based on the proportion of questions answered correctly (Deutsch et al., 2021a). Finally BLEURT is a learned MT metric which predicts a translation quality score using encoded BERT representations of the reference and candidate translations (Sellam et al., 2020).

Implementation details can be found in Appendix A.

4 Metric Optimization

Since reference-free metrics are equivalent to models, then it is possible to define inference procedures which produce the best-possible outputs according to the metrics. Here, we discuss three such (approximate) inference procedures. Importantly, they can all be run at test time because they do not rely on a reference text.

4.1 Direct Optimization

If a reference-free metric scores a candidate output in a way that an efficient approximate inference procedure can be defined, then finding the best possible output under the metric is straightforward.

Among the metrics analyzed in this paper, only Prism-src falls into this category. Because Prism-src assigns a score to a translation equal to its average log-probability under a learned sequence-to-sequence MT model, the approximate best translation under Prism-src can be found by running beam search with the MT model conditioned on the source text.

4.2 Greedy Optimization for Extractive Summarization

Summarization models are generally categorized as being either extractive or abstractive. Extractive systems create a summary by selecting k salient document sentences, whereas abstractive systems typically autoregressively generate a summary with a sequence-to-sequence model.

The best possible extractive summary according to a reference-free metric can be found by enumerating all possible summaries of k sentences, scoring them with the metric, and selecting the summary with the highest score. Since the number of k sentence summaries may be large, this may be computationally expensive. However, an approximate inference procedure can be used instead.

Rather than enumerate all possible extractive summaries, the approximate inference algorithm constructs a summary by greedily selecting one sentence that increases the score of the metric the most (Lin and Bilmes, 2011). This is repeated until a target summary length of k sentences is reached, resulting in an approximation of the best possible summary under the reference-free metric.

A near-identical procedure is commonly used for creating sentence-level labels for training extractive summarization models, except a reference-based evaluation metric, such as ROUGE, is typically used for scoring the sentences instead of a reference-free metric (Nallapati et al., 2017). The key difference is that the output summary from the reference-based procedure is used to train a model which later predicts k salient sentences during inference, whereas the reference-free procedure can be directly used during inference (i.e., without training) to pick the approximately best summary under the reference-free metric.

4.3 Reranking

Exact inference for any reference-free metric can be performed by enumerating all possible outputs, calculating the score of each one, and selecting the output with the highest score. However, it is almost certainly true that this is computationally intractable for any practical application of text generation due to the size of the output space.

To that end, we propose to use reranking (Shen et al., 2004; Och et al., 2004) as an approximate inference procedure in which a pre-trained model for the task at hand is used to restrict the search space to a small set of high-quality candidate out-

puts. These outputs are then scored and reranked using the reference-free metric to identify an approximately best output under the metric.

In practice, we identify a set of k high-quality outputs using standard beam search with pre-trained sequence-to-sequence summarization and MT models and a beam size of k . The top- k partial outputs sorted by their log-likelihood under the pre-trained models are kept at each step of beam search. The final outputs are then reranked by a reference-free metric. For summarization, we use BART (Lewis et al., 2020) trained on the CNN/DailyMail dataset. For MT, we use Facebook’s submission to the WMT’19 translation shared task (Ng et al., 2019). The model is available for en→de, de→en, en→ru, and ru→en.

5 Analysis

5.1 Approximate Inference Effectiveness

Although inference methods for the reference-free metrics can be defined, it is possible that they fail to find high-scoring outputs due to the complexity of the search problem. However in this analysis, we show that the simple approximate inference procedures defined in §4 are effective at optimizing the metrics’ scores.

We compared the outputs obtained by the inference algorithms to those from systems included in the WMT’19, SummEval, and REALSumm datasets. Fig. 2 evaluates using the direct optimization procedure (§4.1) to select the best Prism-src output, Fig. 3 shows the results of using reranking (§4.3) to pick the best outputs according to COMET-QE, and Fig. 4 contains the results of using the greedy extractive procedure (§4.2) to optimize QuestEval. The Figures also include the systems’ scores under the reference-based metrics BLEURT for MT and ROUGE for summarization. Other combinations of reference-based metrics and inference algorithms can be found in Appendix B.

In all MT language pairs and both summarization datasets, the inference algorithms produce the highest scoring outputs under the reference-free metrics, often by a large margin. For example, reranking translations according to their COMET-QE scores on de→en results in a relative 38% improvement in COMET-QE over the WMT’19 submission with the best COMET-QE score (0.347 → 0.478). Clearly, the simple inference procedures can be used to find very high scoring outputs under the reference-free metrics even if the metric does

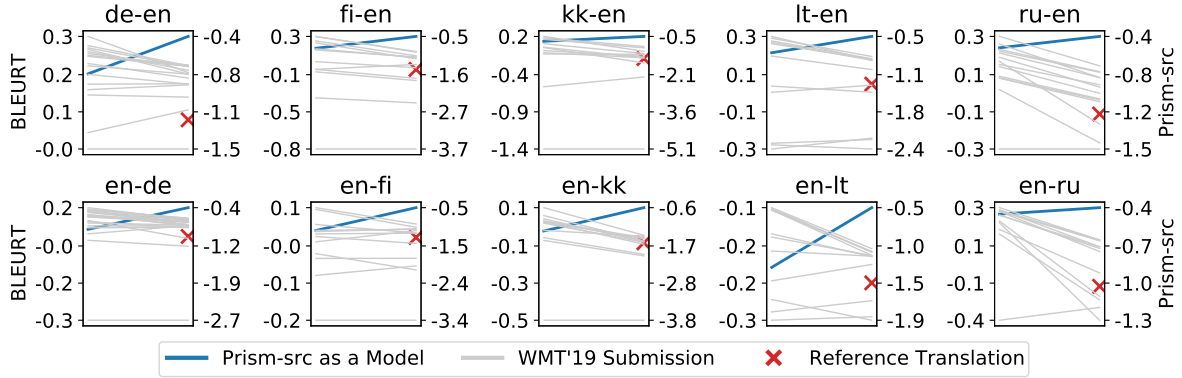


Figure 2: Each line in these plots corresponds to one system evaluated under two different metrics on either y-axis. This illustrates the change in system ranking between the two metrics. We see that directly optimizing Prism-src (blue line; §4.1) yields the highest Prism-src performance (right y-axis) but only an average system as evaluated by BLEURT (left y-axis). The reference translation (red “x”) has a lower Prism-src score compared to many systems across all language pairs, demonstrating Prism-src’s biases toward learned model output and against human-written translations.

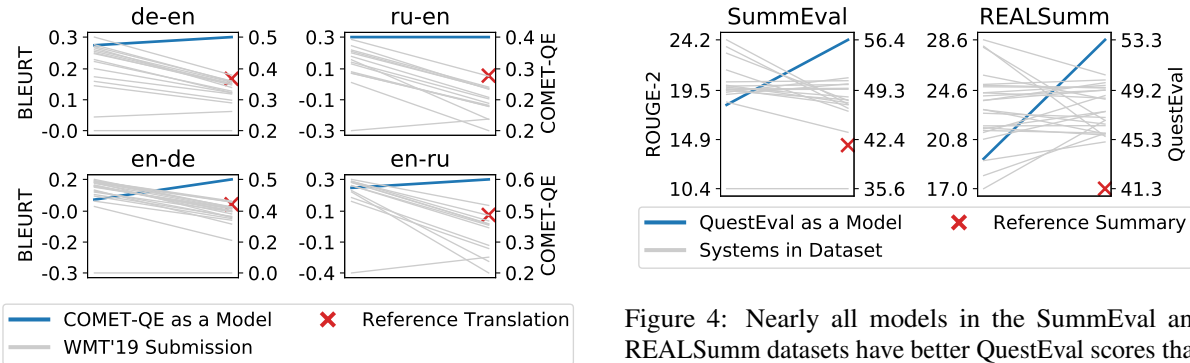


Figure 3: Reranking the output of a pre-trained model results in COMET-src scores which are far higher than the reference translations’ scores (red “x”), demonstrating that a better COMET-src value means the translation is more similar to the metric’s underlying model instead of human-written text.

not directly support efficient inference.

Despite the improvements in reference-free scores, it does not appear as if these outputs are as high-quality as the reference-free metric scores would indicate. Ideally, the outputs would be rated by humans to establish a ground-truth quality score, but a fair comparison to the other systems’ outputs included in the datasets would require re-judging their translations or summaries, which is prohibitively expensive. Instead, we use reference-based metrics as indicators of quality.

When the outputs from our inference algorithms are compared to other systems using reference-based metrics (also shown in Figs. 2, 3, and 4), we

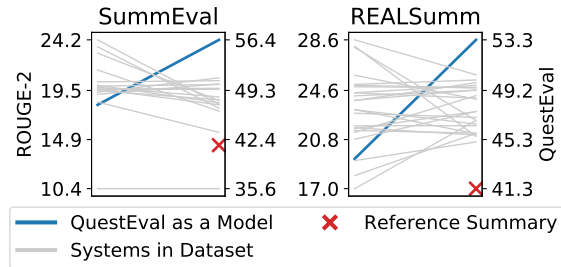


Figure 4: Nearly all models in the SummEval and REALSumm datasets have better QuestEval scores than the reference summaries, demonstrating the metrics’ bias toward learned model output over human-written text.

see that they are often of average quality or worse. For example, the greedy extractive summaries obtained by optimizing QuestEval on REALSumm are among the lowest-performing in the dataset according to ROUGE-2 (Fig. 4). Thus, directly optimizing the reference-free metrics does not always yield a high-quality system, at least according to reference-based metrics (examples shown in Fig 1).

5.2 Undesirable Metric Biases

Ideally, evaluation metrics should score human-written text higher than learned model outputs since it is very likely that the human references are of higher-quality. However, we see that this does not always happen with reference-free metrics.

Figs. 2, 3, and 4 additionally contain the scores of the reference texts under the reference-free metrics (marked with a red “x”). In all settings, the

inference algorithms’ outputs score higher than the references. This is unsurprising because they select their outputs to optimize the metrics’ values. However, it demonstrates that as models continually improve their reference-free scores, they will begin to converge onto the metrics’ underlying models and not onto human-quality text. As the research goal of text generation is to build models which produce human-like text, the reference-free metrics’ scores do not align with this goal.

Further, the same figures show that it is often the case that other models in the datasets—which did not directly optimize the reference-free metrics—also are scored higher than human text by the reference-free metrics. This is especially true for Prism-src and QuestEval, but less so for COMET-QE. For example, on language pair de-en, only one system has a lower Prism-src score than the reference translations (Fig. 2). These metrics appear to have a bias for outputs from learned models over human-written text.

We suspect this bias for outputs from learned models is due to how the internal models used by Prism-src and QuestEval to score text are trained. Prism-src scores translations using an MT model that was trained on standard MT data. QuestEval uses a question-weighting model that predicts how likely a question is answered in a reference summary, which was trained on the same CNN/DailyMail dataset that the summarization models in the summarization datasets were also trained on. Thus, these metrics internally use models which are directly or indirectly trained to perform the generation task (MT or summarization).

Generation models which are trained on the same datasets are known to exhibit similar behavior and make similar mistakes, and their outputs often look markedly different from human-written text. Therefore, we suspect that the signals the internal MT/question-weighting models have learned to identify high-quality text are similar to those which the task-specific models have learned to produce their output, and thus receive high scores by the metrics. In contrast, the human-written text likely does not rely on these signals, and is thus perceived as low-quality by the metrics.

This bias toward learned model outputs is potentially less severe for COMET-QE because unlike Prism-src and QuestEval, it is specifically trained to predict translation quality using manually collected human judgments. Those human judgments also

contain evaluations of human-written translations, so it is better at distinguishing human-versus-model output (Freitag et al., 2021). Thus, the signals it learns to identify high-quality text are likely different than what is learned by the translation models in the WMT’19 dataset.

In summary, the metrics’ biases toward their own outputs (or other learned model outputs) and against human texts demonstrates they reward outputs which look more like their own instead of human-quality text, an undesirable property of an evaluation metric.

5.3 Reference-Free Metrics as Pseudo-References

Thus far, we have argued that the underlying model of a reference-free metric is the theoretical best model according to the metric. It would intuitively follow that the more similar another model is to the metric’s underlying model, the higher metric score that model would receive. To that end, in this analysis we demonstrate that scoring a system with a reference-free metric is roughly equivalent to evaluating that system’s outputs against the metric’s best possible output using a reference-based metric. This further demonstrates the limitations of reference-free metrics, including their biases toward their own underlying models’ outputs.

A pseudo-reference is a piece of text which is used in place of a human-written reference to evaluate some candidate text with a reference-based metric (Louis and Nenkova, 2013; Gao et al., 2020). For the reference-free metrics, we define the pseudo-reference to be the output from the inference procedures defined in §4 (i.e., those evaluated in §5.1). For example, the QuestEval pseudo-reference is the extractive summary which was selected to greedily maximize the QuestEval score.

Once the pseudo-references have been defined, they can be used in conjunction with any reference-based metric, such as BLEURT or QAEval, to evaluate other translations or summaries. To quantify the similarity between evaluating a system with a pseudo-reference and a reference-free metric, we calculated the Pearson correlation between the system-level scores between the two methods. These correlations are shown in Figs. 5, 6, and 7.

For MT, we find that the BERTScore of a translation that uses pseudo-references instead of the human-written reference has an average Pearson correlation of 0.95 and 0.92 to the Prism-src (Fig. 5)

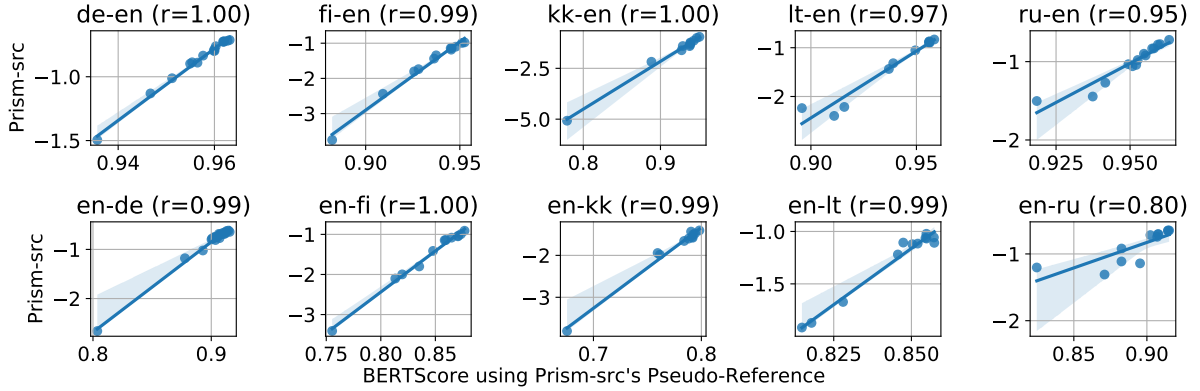


Figure 5: A system’s BERTScore in which the output from directly optimizing Prism-src is used as the reference (x-axis) is strongly correlated to that same system’s Prism-src score (y-axis). This demonstrates Prism-src is roughly equivalent to evaluating systems with a pseudo-reference translation which is generated by a model. Pearson’s r shown in the title of each plot.

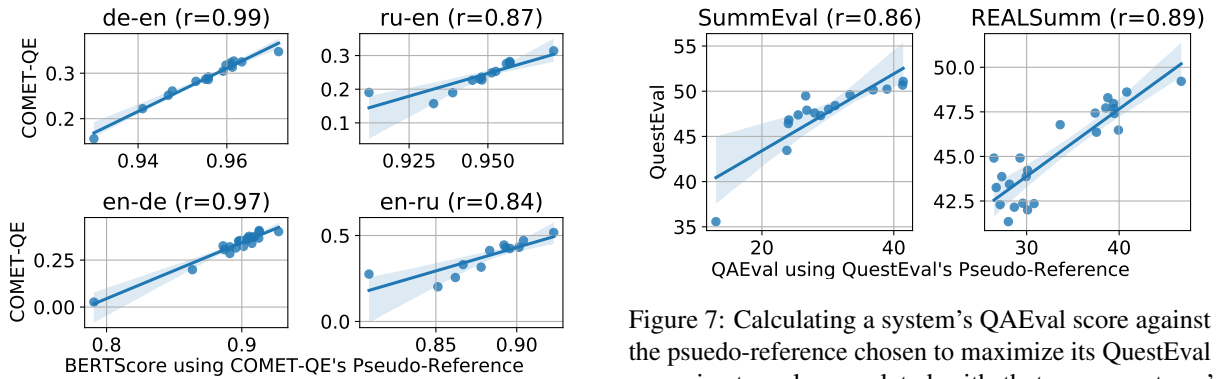


Figure 6: Scoring systems with BERTScore against psudeo-references obtained by optimizing COMET-src strongly correlates to the systems’ COMET-src scores.

and COMET-QE scores (Fig. 6), respectively. The summarization correlations for QuestEval to QA-Eval using a pseudo-reference (Fig. 7) are also rather strong at 0.88 on average. The correlations using other reference-based metrics are slightly weaker on average due to low values on specific datasets or language pairs (see Appendix B), but there are many instances in which the correlations are ≥ 0.9 .

Overall, these correlations are very strong, suggesting that the reference-free metrics are roughly equivalent to using pseudo-references to evaluate other models. Once the metrics are viewed this way, their limitations become clear. The metrics’ outputs are the gold-standard against which all other outputs should be evaluated. Thus, the metrics favor their own outputs (the pseudo-references) and other outputs which are more similar (where similarity is measured using reference-based metrics).




Figure 7: Calculating a system’s QA-Eval score against the psudeo-reference chosen to maximize its QuestEval score is strongly correlated with that same systems’ QuestEval score on SummEval and REALSumm.

Further, their ability to evaluate other models is inherently limited by the qualities of their pseudo-references. If a system outputs a translation or a summary which is higher-quality than the pseudo-reference, it will be incorrectly penalized because it is different than the pseudo-reference even though those differences are actually improvements. Thus, the metrics’ scores of systems which are better in quality than their own underlying models will be misleading.

6 Discussion

6.1 Reference-Free Evaluation

Due to the limitations of reference-free metrics that we demonstrated in previous Sections, we argue that reference-free metrics should not be used to measure progress on a task, for instance, by concluding that an MT or summarization model is better than another because it has a higher reference-free metric score. If they are used this way, the

model which will achieve the best performance is already known (the metric’s underlying model; §2), and simple inference procedures are effective at using those models to generate high-scoring outputs (§4, §5.1). Improving the metric’s score is no longer about creating better models on the training data for the task; rather, it is about coming up with better procedures to optimize the metric during inference. In the end, the quality of the system which is produced by this evaluation methodology is limited by the quality of the metric itself.

When references are not available, we recommend investing resources into collecting human-written references that can be used for evaluation purposes instead of pursuing reference-free metrics. Although individual reference-based metrics have their own flaws (Graham, 2015; Bhandari et al., 2020; Deutsch et al., 2021b), the class of reference-based metrics still ultimately encourages systems to generate text which is similar to humans, which is the goal of text generation research as it is currently defined.

6.2 What about Inherently Reference-Free Evaluations?

Although we argue that measuring system quality with reference-free metrics is flawed and misleading, abandoning reference-free evaluations completely is not an option and one which we do not advocate for. There are many aspects of generated text that inherently do not rely on the presence of a reference to be evaluated, such as the fluency of a translation or how faithful a summary is to its input document. There is no obvious benefit of including a reference text in these evaluations. As such, they suffer from the same issues as the metrics analyzed in this work, yet the motivation for being able to automatically evaluate these aspects of text is clear.

In these scenarios, we recommend using reference-free metrics as diagnostic tools for better understanding the behavior of models instead of a method for measuring progress on a task. For instance, the perplexity of a summary under a large-scale language model is a useful statistic to report in order to approximate its fluency, but the value should only be interpreted as exactly what it measures—how likely the observed text is under the language model—with the understanding that the measure is inherently biased towards the underlying language model. The perplexity should not be used to drive research on how to generate more

fluent summaries because the most fluent summarization model is the language model itself.

This recommendation applies not only to metrics which measure inherently reference-free aspects of text, but also to the metrics that evaluate aspects which we argue should use references, such as those analyzed in this work. They are certainly useful statistics to report, but improving their values as much as possible should not be the goal.

7 Related Work

Various other reference-free metrics have been proposed for MT (Fonseca et al., 2019; Rei et al., 2021), summarization (Louis and Nenkova, 2013; Scialom et al., 2019; Xenouelas et al., 2019; Vasilyev et al., 2020; Scialom et al., 2021), dialog generation (Mehri and Eskenazi, 2020; Honovich et al., 2021), image captioning (Hessel et al., 2021), and simplification (Martin et al., 2018; Kriz et al., 2020).

Reference-free metrics for MT were evaluated and compared to reference-based metrics in the WMT’21 metrics shared task (Freitag et al., 2021). In some evaluation settings, reference-free metrics perform better than reference-based metrics. Durmus et al. (2022) demonstrate that reference-free metrics evaluate text using undesirable, spurious correlations and propose methods for avoiding learning from those signals.

While many of the work which propose reference-free metrics do not explicitly state that they are trying to replace reference-based metrics, we worry that this goal is implied and that the limitations of their metrics are not stated clearly enough. For example, simply by arguing that their reference-free metrics have stronger correlations to human judgments than their reference-based counterparts, an uninformed reader may conclude that the two types of metrics are interchangeable and use a reference-free metric instead of a reference-based metric without fully understanding the consequences. We recommend authors be explicit about the limitations of reference-free metrics and say, for example, they can be used to complement existing reference-based evaluations rather than a replacement (Louis and Nenkova, 2013).

One such application of reference-free metrics which fits our recommendations for how they should be used is the WMT shared task on quality estimation (QE; Specia et al., 2020, 2021). QE metrics are used to flag translations which might

require post-editing or to identify “catastrophic errors” in translations, for instance hateful or violent text that was hallucinated by the model, without the aid of a reference. The intended purpose of QE metrics is not to rank and compare MT systems. We argue that the “rebranding” of QE metrics as reference-free metrics by some authors contributes to the notion that they can and should be used the same way that reference-based metrics are and that they are inherently better because they do not require a reference. We recommend that future proposals of methods for evaluating text without references are called QE metrics instead of reference-free metrics to make the distinction more clear.

In previous work, Prism-src was further explored by [Agrawal et al. \(2021\)](#), who experimented with its model capacity, scoring methods, and more. They perform an analysis using pseudo-references similar to ours in §5.3. In their experiment, they show that the cross-BLEU score calculated between Prism-src’s pseudo-reference and the system output does not appear to be correlated to Prism-src’s system ranking relative to human-based system ranking. That is, a higher cross-BLEU score does not necessarily result in Prism-src ranking a system higher than it should. They conclude that their analysis does not show evidence that Prism-src is biased toward outputs which are more similar to its own, which seemingly contradicts our own experiments and results in §5.3. However, we argue that the high correlations between a system’s Prism-src score and its similarity to the metric’s pseudo-reference (measured by reference-free metrics) does demonstrate such a bias exists, but [Agrawal et al. \(2021\)](#) did not find evidence this bias negatively impacted the rankings of the systems.

Our experiments in §4.3 leverage reranking to optimize reference-free metrics’ scores. Reranking in machine translation has a long history ([Shen et al., 2004](#); [Och et al., 2004](#)) and has been done with reference-free metrics, for instance in concurrent work by [Fern et al. \(2022\)](#). The aim of our experiment was not to propose a novel technique, but rather to point out that reference-free metrics can be optimized in this way, and doing so may be potentially detrimental. Reference-based metrics can also be used for reranking, for example, by using Minimum Bayes Risk (MBR) decoding ([Kumar and Byrne, 2004](#); [Freitag et al., 2022](#); [Fern et al., 2022](#)), which can bias the output toward a particular metric. While techniques like MBR ap-

proximate the reference-based reward function, if a system is evaluated with a reference-free metric, the exact reference-free reward function can be directly optimized during inference.

8 Conclusion

In this work, we have argued that reference-free metrics are inherently limited in their ability to evaluate generated text. Because they are equivalent to generation models, they can be directly optimized at test time, they are biased toward their own models’ outputs and outputs from similar models, and they can be biased against higher-quality text, such as human-written references. Therefore, we recommend against using reference-free metrics as measures of progress on tasks and instead advocate for them to be used as useful statistics to calculate in order to better understand model behavior.

Acknowledgments

The authors would like to thank the ACL ARR reviewers, Markus Freitag, and Colin Cherry for their insightful feedback on our work. Their comments raised many interesting points and have helped improved the quality of our final publication.

This work was supported by Contracts FA8750-19-2-1004 and FA8750-19-2-0201 with the US Defense Advanced Research Projects Agency (DARPA). Approved for Public Release, Distribution Unlimited. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA Contract No. 2019-19051600006 under the BETTER Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, the Department of Defense, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

This research is supported by a Focused Award from Google.

The second author is supported by the Eric and Wendy Schmidt Postdoctoral Award for Women in Mathematical and Computing Sciences.

Limitations

Our work argues that reference-free metrics are theoretically biased toward models which are similar to the metrics' underlying models and biased against human-written output, and we show experimental evidence to support this argument. While our argument applies to the class of reference-free metrics regardless of the task, our experiments use only one translation and two summarization datasets, so the extent to which this problem can be experimentally demonstrated on other tasks is not known. However, we suspect it is possible.

Then, although we point out the limitations and biases of reference-free metrics, we do not provide a clear alternative for how to automatically evaluate inherently reference free aspects of text that would not have these problems. It is not clear if this is possible or how to proceed.

Our analysis in §5.1 uses reference-based metrics as surrogates for human judgments to determine system output quality. Ideally we would use humans to determine the true quality of the outputs from our inference algorithms, but that is out of the scope for this paper.

References

- Sweta Agrawal, George Foster, Markus Freitag, and Colin Cherry. 2021. [Assessing Reference-Free Peer Evaluation for Machine Translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1158–1171, Online. Association for Computational Linguistics.
- Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. [Re-evaluating Evaluation in Text Summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9347–9359, Online. Association for Computational Linguistics.
- Daniel Deutsch, Tania Bedrax-Weiss, and Dan Roth. 2021a. [Towards Question-Answering as an Automatic Metric for Evaluating the Content Quality of a Summary](#). *Transactions of the Association for Computational Linguistics*, 9:774–789.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2021b. [A Statistical Analysis of Summarization Evaluation Metrics Using Resampling Methods](#). *Transactions of the Association for Computational Linguistics*, 9:1132–1146.
- Daniel Deutsch and Dan Roth. 2020. [SacreROUGE: An Open-Source Library for Using and Developing Summarization Evaluation Metrics](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 120–125, Online. Association for Computational Linguistics.
- Daniel Deutsch and Dan Roth. 2022. [Repro: An Open-Source Library for Improving the Reproducibility and Usability of Publicly Available Research Code](#). *ArXiv*, abs/2204.13848.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Esin Durmus, Faisal Ladhak, and Tatsunori Hashimoto. 2022. [Spurious Correlations in Reference-Free Evaluation of Text Generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1443–1454, Dublin, Ireland. Association for Computational Linguistics.
- Alexander Fabbri, Wojciech Kryscinski, Bryan McCann, R. Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating Summarization Evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Fern, Patrick es, António Farinhas, Ricardo Rei, José De Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. [Quality-Aware Decoding for Neural Machine Translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.
- Erick Fonseca, Lisa Yankovskaya, André F. T. Martins, Mark Fishel, and Christian Federmann. 2019. [Findings of the WMT 2019 Shared Tasks on Quality Estimation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–12, Florence, Italy. Association for Computational Linguistics.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. [High Quality Rather than High Model Probability: Minimum Bayes Risk Decoding with Neural Metrics](#). *Transactions of the Association for Computational Linguistics*, 10:811–825.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021. [Results of the WMT21 Metrics Shared Task: Evaluating Metrics with Expert-based Human Evaluations on TED and News Domain](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.

- Yang Gao, Wei Zhao, and Steffen Eger. 2020. **SUPERT: Towards New Frontiers in Unsupervised Evaluation Metrics for Multi-Document Summarization**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1347–1354, Online. Association for Computational Linguistics.
- Yvette Graham. 2015. **Re-evaluating Automatic Summarization with BLEU and 192 Shades of ROUGE**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 128–137, Lisbon, Portugal. Association for Computational Linguistics.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Roman Le Bras, and Yejin Choi. 2021. **CLIPScore: A Reference-free Evaluation Metric for Image Captioning**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Or Honovich, Leshem Choshen, Roei Aharoni, Ella Neeman, Idan Szpektor, and Omri Abend. 2021. **Q²: Evaluating Factual Consistency in Knowledge-Grounded Dialogues via Question Generation and Question Answering**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7856–7870, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Reno Kriz, Marianna Apidianaki, and Chris Callison-Burch. 2020. **Simple-QE: Better Automatic Quality Estimation for Text Simplification**. *arXiv preprint arXiv:2012.12382*.
- Shankar Kumar and William Byrne. 2004. **Minimum Bayes-Risk Decoding for Statistical Machine Translation**. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 169–176, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising Sequence-to-Sequence Pretraining 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.
- Chin-Yew Lin. 2004. **ROUGE: A Package for Automatic Evaluation of Summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Hui Lin and Jeff Bilmes. 2011. **A Class of Submodular Functions for Document Summarization**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 510–520, Portland, Oregon, USA. 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**. *CoRR*, abs/1907.11692.
- Annie Louis and Ani Nenkova. 2013. **Automatically Assessing Machine Summary Content Without a Gold Standard**. *Computational Linguistics*, 39(2):267–300.
- Qingsong Ma, Johnny Wei, Ondřej Bojar, and Yvette Graham. 2019. **Results of the WMT19 Metrics Shared Task: Segment-Level and Strong MT Systems Pose Big Challenges**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 62–90, Florence, Italy. Association for Computational Linguistics.
- Louis Martin, Samuel Humeau, Pierre-Emmanuel Mazaré, Éric de La Clergerie, Antoine Bordes, and Benoît Sagot. 2018. **Reference-less Quality Estimation of Text Simplification Systems**. In *Proceedings of the 1st Workshop on Automatic Text Adaptation (ATA)*, pages 29–38, Tilburg, the Netherlands. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2020. **USR: An Unsupervised and Reference Free Evaluation Metric for Dialog Generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. **SummaRuNNer: A Recurrent Neural Network Based Sequence Model for Extractive Summarization of Documents**. In *Proc. of the Conference on Artificial Intelligence (AAAI)*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. **Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond**. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. **Facebook FAIR’s WMT19 News Translation Task Submission**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 314–319, Florence, Italy. Association for Computational Linguistics.
- Franz Josef Och, Daniel Gildea, Sanjeev Khudanpur, Anoop Sarkar, Kenji Yamada, Alex Fraser, Shankar Kumar, Libin Shen, David Smith, Katherine Eng, Viren Jain, Zhen Jin, and Dragomir Radev. 2004.

- A Smorgasbord of Features for Statistical Machine Translation. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 161–168, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **BLEU: a Method for Automatic Evaluation of Machine Translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. **A Call for Clarity in Reporting BLEU Scores**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Ricardo Rei, Ana C Farinha, Chrysoula Zerva, Daan van Stigt, Craig Stewart, Pedro Ramos, Taisiya Glushkova, André FT Martins, and Alon Lavie. 2021. **Are References Really Needed? Unbabel-IST 2021 Submission for the Metrics Shared Task**. In *Proceedings of the Sixth Conference on Machine Translation, Online*. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. **COMET: A Neural Framework for MT Evaluation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. **QuestEval: Summarization Asks for Fact-based Evaluation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. **Answers Unite! Unsupervised Metrics for Reinforced Summarization Models**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3246–3256, Hong Kong, China. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. **BLEURT: Learning Robust Metrics for Text Generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Libin Shen, Anoop Sarkar, and Franz Josef Och. 2004. **Discriminative Reranking for Machine Translation**. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 177–184, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. **Findings of the WMT 2020 Shared Task on Quality Estimation**. In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Chrysoula Zerva, Zhenhao Li, Vishrav Chaudhary, and André F. T. Martins. 2021. **Findings of the WMT 2021 Shared Task on Quality Estimation**. In *Proceedings of the Sixth Conference on Machine Translation*, pages 684–725, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020. **Automatic Machine Translation Evaluation in Many Languages via Zero-Shot Paraphrasing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Oleg Vasilyev, Vedant Dharnidharka, and John Bohannon. 2020. **Fill in the BLANC: Human-free quality estimation of document summaries**. In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 11–20, Online. Association for Computational Linguistics.
- Stratos Xenoules, Prodromos Malakasiotis, Marianna Apidianaki, and Ion Androutsopoulos. 2019. **SUMQE: a BERT-based Summary Quality Estimation Model**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6005–6011, Hong Kong, China. Association for Computational Linguistics.
- Shiyue Zhang and Mohit Bansal. 2021. **Finding a Balanced Degree of Automation for Summary Evaluation**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6617–6632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. **BERTScore: Evaluating Text Generation with BERT**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A Implementation Details

Our experiments were run on a single Titan RTX GPU with 24 GB of memory. Since the experi-

ments did not require training models and are limited to using metrics to score model outputs, they were relatively inexpensive; most experiments completed in several hours.

The ROUGE and BLEU implementations are from SacreROUGE (Deutsch and Roth, 2020) and SacreBLEU (Post, 2018), respectively. We also used SacreROUGE for QAEval. For BLEURT, we use the BLEURT-base-128 version. BERTScore uses the embeddings from RoBERTa-Large for English and the default contextual embeddings for the non-English languages. (Liu et al., 2019). For QuestEval, we used version v0.1.1 of the corresponding library⁴ based on the authors’ recommendation in their documentation. The Prism-src and COMET-QE implementations were based on the authors’ original code. All of the metrics implementations can be found in the Repro library (Deutsch and Roth, 2022).

B Additional Results

This Section contains additional results to supplement the main body of the paper.

Table 1 contains the Figure numbers for additional combinations of reference-free metrics, inference algorithms, and reference-based metrics to what was presented in §5.1. In every setup, the inference procedure generated outputs which performed the best compared to all of the other available systems in the datasets.

Table 2 contains the correlations calculated between the systems’ Prism-src scores and their reference-based scores in which the pseudo-reference is determined by running an inference algorithm on Prism-src. This experiment was explained in §5.3. Tables 3 and 4 contain the same results but for COMET-QE and QuestEval.

⁴<https://github.com/ThomasScialom/QuestEval>

Reference-Free Metric	Optimization Method	Reference-Based Metric	Figure
Prism-src	Direct Optimization (§4.1)	BLEURT	Fig. 8
Prism-src	Direct Optimization (§4.1)	BLEU	Fig. 9
Prism-src	Direct Optimization (§4.1)	BERTScore	Fig. 10
Prism-src	Reranking (§4.3)	BLEURT	Fig. 11
Prism-src	Reranking (§4.3)	BLEU	Fig. 12
Prism-src	Reranking (§4.3)	BERTScore	Fig. 13
COMET-QE	Reranking (§4.3)	BLEURT	Fig. 3
COMET-QE	Reranking (§4.3)	BLEU	Fig. 14
COMET-QE	Reranking (§4.3)	BERTScore	Fig. 15
QuestEval	Greedy Extractive (§4.2)	ROUGE	Fig. 4
QuestEval	Greedy Extractive (§4.2)	BERTScore	Fig. 16
QuestEval	Greedy Extractive (§4.2)	QAEval	Fig. 17
QuestEval	Reranking (§4.3)	ROUGE	Fig. 18
QuestEval	Reranking (§4.3)	BERTScore	Fig. 19
QuestEval	Reranking (§4.3)	QAEval	Fig. 20

Table 1: The Figure on each row corresponds to using the optimization method (§4) to generate an output with the highest reference-free metric score, which is then evaluated with the reference-free and reference-based metric. In each Figure, the systems’ reference-based scores are on the left y-axis, the reference-free score on the right y-axis, the outputs which were generated by running inference on the reference-free metric in blue, and the reference text’s reference-free metric score in the red “x.”

Metric	de-cs	de-en	de-fr	en-cs	en-de	en-fi	en-kk	en-lt	en-ru	en-zh	fi-en	fr-de	kk-en	lt-en	ru-en	zh-en	Avg.
<i>Direct Optimization (§4.1)</i>																	
BLEURT	0.99	0.97	0.69	0.92	0.99	0.98	0.99	0.93	0.61	0.73	0.99	0.16	0.98	0.99	0.84	0.99	0.86
BLEU	0.95	0.92	0.62	0.94	0.96	0.93	0.96	0.91	0.77	0.55	0.98	0.47	0.86	0.97	0.82	0.98	0.85
BERTScore	0.98	1.00	0.90	1.00	0.99	1.00	0.99	0.99	0.80	0.87	0.99	0.82	1.00	0.97	0.95	0.99	0.95
<i>Reranking (§4.3)</i>																	
BLEURT	-	0.92	-	-	0.94	-	-	-	0.66	-	-	-	-	-	0.85	-	0.84
BLEU	-	0.83	-	-	0.89	-	-	-	0.82	-	-	-	-	-	0.80	-	0.83
BERTScore	-	0.95	-	-	0.97	-	-	-	0.83	-	-	-	-	-	0.95	-	0.92

Table 2: The correlation between (1) the Prism-src score for the systems submitted to WMT’19 and (2) the same systems’ reference-based metric (left column) score calculated against the pseudo-reference for Prism-src that was generated using the direct optimization inference algorithm (top) or reranking (bottom).

Metric	de-en	en-de	en-ru	ru-en	Avg.
BLEURT	0.98	0.96	0.61	0.76	0.83
BLEU	0.91	0.95	0.86	0.80	0.88
BERTScore	0.99	0.97	0.84	0.87	0.92

Table 3: The correlation between (1) the COMET-QE score for the systems submitted to WMT’19 and (2) the same systems’ reference-based metric (left column) score calculated using the pseudo-reference for COMET-QE that was generated using the reranking inference algorithm.

Metric	SummEval	REALSumm	Avg.
<i>Greedy Extractive (§4.2)</i>			
ROUGE-2	0.85	0.83	0.84
BERTScore	0.62	0.85	0.74
QAEval	0.86	0.89	0.88
<i>Reranking (§4.3)</i>			
ROUGE-2	0.77	0.77	0.77
BERTScore	0.19	0.78	0.49
QAEval	0.87	0.85	0.86

Table 4: The correlation between (1) the QuestEval score for the systems included in the REALSumm and SummEval datasets and (2) the same systems’ reference-based metric (left column) score calculated against the pseudo-reference for QuestEval that was generated using the greedy extractive summarization inference algorithm (top) or reranking (bottom).




Figure 8: See Table 1 for a description of this Figure.

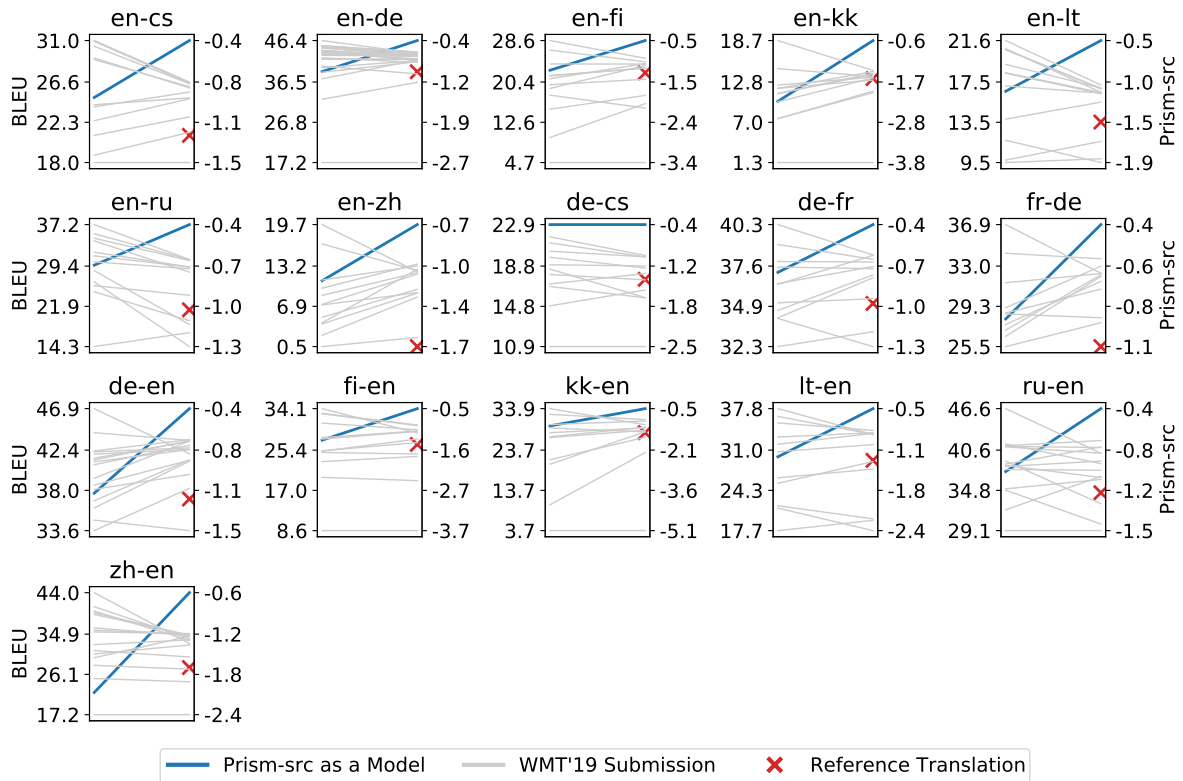


Figure 9: See Table 1 for a description of this Figure.




Figure 10: See Table 1 for a description of this Figure.




Figure 11: See Table 1 for a description of this Figure.

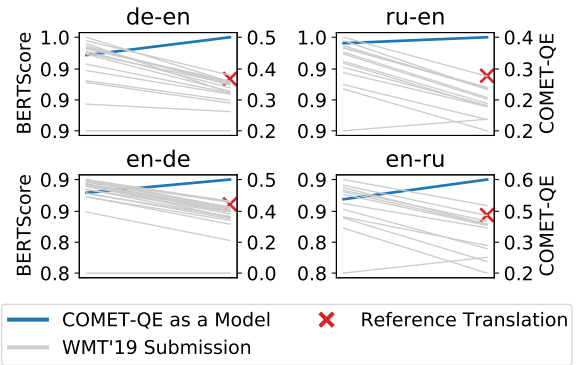


Figure 15: See Table 1 for a description of this Figure.




Figure 12: See Table 1 for a description of this Figure.




Figure 16: See Table 1 for a description of this Figure.

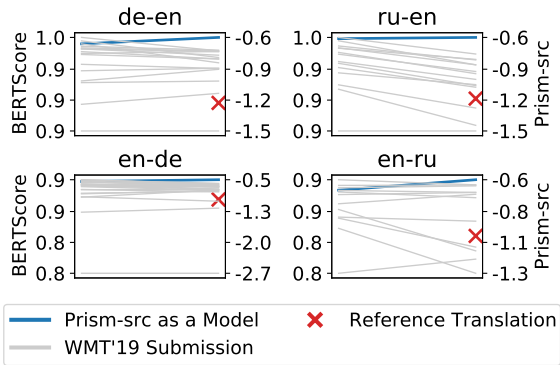


Figure 13: See Table 1 for a description of this Figure.




Figure 17: See Table 1 for a description of this Figure.




Figure 14: See Table 1 for a description of this Figure.




Figure 18: See Table 1 for a description of this Figure.




Figure 19: See Table 1 for a description of this Figure.




Figure 20: See Table 1 for a description of this Figure.