

# Sentiment-Aware Measure (SAM) for Evaluating Sentiment Transfer by Machine Translation Systems

**Hadeel Saadany**

RGCL  
University of Wolverhampton  
Wolverhampton, UK  
h.s.saadany@wlv.ac.uk

**Constantin Orăsan**

Centre for Translation Studies  
University of Surrey  
Guildford, UK  
c.orasan@surrey.ac.uk

**Emad Mohamed**

RGCL  
University of Wolverhampton  
Wolverhampton, UK  
e.mohamed2@wlv.ac.uk

**Ashraf Tantavy**

School of Computer Science and Informatics  
De Montfort University  
Leicester, UK  
ashraf.tantavy@dmu.ac.uk

## Abstract

In translating text where sentiment is the main message, human translators give particular attention to sentiment-carrying words. The reason is that an incorrect translation of such words would miss the fundamental aspect of the source text, i.e. the author's sentiment. In the online world, MT systems are extensively used to translate User-Generated Content (UGC) such as reviews, tweets, and social media posts, where the main message is often the author's positive or negative attitude towards the topic of the text. It is important in such scenarios to accurately measure how far an MT system can be a reliable real-life utility in transferring the correct affect message. This paper tackles an under-recognised problem in the field of machine translation evaluation which is judging to what extent automatic metrics concur with the gold standard of human evaluation for a correct translation of sentiment. We evaluate the efficacy of conventional quality metrics in spotting a mistranslation of sentiment, especially when it is the sole error in the MT output. We propose a numerical 'sentiment-closeness' measure appropriate for assessing the accuracy of a translated affect message in UGC text by an MT system. We will show that incorporating this sentiment-aware measure can significantly enhance the correlation of some available quality metrics with the human judgement of an accurate translation of sentiment.

## 1 Introduction

Standard quality measures for assessing the performance of machine translation systems, such as BLEU (Papineni et al., 2002), are domain agnos-

tic; they evaluate the translation accuracy regardless of the semantic domain or linguistic peculiarities of the source text. Consequently, they give equal penalty weight to inaccurate translation of n-grams, which may lead to performance overestimation (or underestimation). For example, the Arabic high-rated Goodreads book review showing the reviewers overall satisfaction of a novel: 'الرواية رهيبه عيها الوحيد الجزء الاخير' (*The novel is great, its only flaw is the last part*) is mistranslated by the available online translation tool which wrongly outputs a negative sentiment (*The novel is terrible, its only flaw is the last part*). Despite the distortion of the affect message, this translation receives an equally high score as a correct translation (*The story is great, its only flaw is the last part*), but which uses *story* instead of *novel*. This is because BLEU mildly penalises the wrong translation as it swaps only one uni-gram ('great') with its opposite ('terrible'). Yet, the mistranslation of this particular uni-gram is a critical error as it is most pivotal in transferring the sentiment, and hence the MT performance is over-estimated.

There have been numerous efforts to address the common pitfalls of n-gram-based evaluation metrics by incorporating semantic and contextual features. However, despite research evidence of its analytical limitations, BLEU is still the *de facto* standard for MT performance evaluation (Mathur et al., 2020a; Reiter, 2018). Moreover, although the introduction of more semantically-oriented metrics showed a better correlation with human judgement, still the estimation of sentiment preservation in UGC has not yet been investigated. This is par-

tially due to the fact that the domain and linguistic style of the WMTs datasets typically used for metric evaluation (e.g. *newstest2020*, Canadian Parliament, Wikipedia, UN corpus) is quite different than the non-standard noisy UGC where sentiment is the main content of a telegraphic message (Ma et al., 2018; Barrault et al., 2019; Mathur et al., 2020b). Assessed over these WMT datasets, some metrics manifest an almost perfect correlation with human evaluation on the segment-level (e.g. WMT20 participating metrics record results of up to 0.97 Pearson correlation on the *newstest2020*) (Mathur et al., 2020b). However, research has also shown that metrics usually report weaker correlation with low human assessment score ranges (Takahashi et al., 2020a,b).

In this paper, we argue that the high correlation of some metrics may not be replicated with a different domain such as sentiment-oriented UGC, specifically when there is a mistranslation of sentiment-critical word(s). For this reason, we propose a ‘sentiment-closeness’ measure that can accommodate for a better evaluation of the MT system’s ability to capture the correct sentiment of the source text.

To present our sentiment quality measure, we first briefly explain in Section 2 why metrics commonly used for MT quality estimation may not always be efficient in assessing sentiment-critical translation errors. In Section 3, we present an experiment to quantitatively assess the divergence of the analysed metrics from a human judgement of a correct/incorrect translation of an affect message. Section 4 presents our proposed solution for fine-tuning quality metrics for a better correlation with human judgement. Section 5 presents the results of incorporating our sentiment-measure in different quality metrics. In Section 6 we conduct an error analysis of our empirical approach and discuss its limitations. Finally, Section 7 presents a conclusion on our conducted experiments.

## 2 Available Metrics

Automatic evaluation metrics usually take the output of an MT system (*hypothesis*) and compare it to one or several translations produced by human translator(s) (*reference*). Based on their matching methods, the most commonly used automatic metrics can be broadly categorised into: surface n-gram matching and embedding matching. Surface n-gram methods work by calculating exact match-

ing, heuristic matching or an edit distance between the aligned n-grams of the reference and hypothesis translation(s). The embedding methods, on the other hand, calculate a similarity score between learned token representations, such as contextual embedding vectors, with or without the aid of external linguistic tools. In the following sections, we briefly explain the methods behind three canonical metrics as representative of each category. We illustrate why the theoretical foundation of each metric may not be optimum for evaluating the translation of sentiment-oriented UGC.

### 2.1 Surface N-gram Matching Metrics

**BLEU** The standard metric for assessing empirical improvement of MT systems is BLEU (Papineni et al., 2002). Simply stated, the objective of BLEU is to compare n-grams of the candidate translation with n-grams of the reference translation and count the number of matches; the more the matches, the better the candidate translation. The final score is calculated using a modified n-gram precision multiplied by a brevity penalty so that a good candidate translation would match the reference translation in length, in word choice, and in word order. The disadvantage of the BLEU metric which is relevant to our present study is that it treats all n-grams equally. Due to its restrictive surface n-gram matching, it does not account for the semantic importance of an n-gram in the context of a text. Accordingly, BLEU would incorrectly give a high score to an MT output if it scores exact match with the reference except for one uni-gram, even if this uni-gram completely changes the sentiment of a text (e.g. ‘terrible’ and ‘great’ as in the Goodreads example above). Online built-in MT tools have been shown to frequently transfer the exact opposite sentiment word for some dialectal expressions in UGC translated into English (Saadany and Orasan, 2020). Therefore, the BLEU evaluation of an MT performance would be misleadingly over-permissive in such cases where only one or two sentiment-critical words are mistranslated.

**METEOR** METEOR (Banerjee and Lavie, 2005) incorporates semantic information as it evaluates translation by calculating either exact match, stem match, or synonymy match. For synonym matching, it utilises WordNet synsets (Pedersen et al., 2004). More recent versions (METEOR 1.5 and METEOR++2.0) apply also importance

weighting by giving smaller weight to function words (Denkowski and Lavie, 2014; Guo and Hu, 2019). The METEOR score ranges from 0 (worst translation) to 1 (best translation). There are two shortcomings to the METEOR metric which do not make it a robust solution for evaluating sentiment transfer. First, the synonym matching is limited to checking whether the two words belong to the same synset in WordNet. However, WordNet synonymy classification is different than regular thesauruses. For example, ‘glad’ is a synset of ‘happy’ and hence considered a synonym, whereas ‘cheerful’ is not a direct synset to ‘happy’ and hence would be considered a mismatch by METEOR. The following two examples illustrate further the limitations of using WordNet synonymy by METEOR:

### Example 1

Scores: [METEOR: 0.46]

- Hypothesis: “The weather is sunny, what a happy day”
- Reference: “The sun is shining, what a cheerful day”

### Example 2

Scores: [METEOR: 0.48]

- Hypothesis: “I’m not sure why, but I feel so happy today”
- Reference: “I don’t get it, but I feel so sad today”

The METEOR scores for Examples 1 and 2 clearly diverge from a human evaluation of a good translation. In the first example, although the translation conveys the correct emotion (‘happiness’), it receives a similar METEOR score to the hypothesis in the second example which gives the exact opposite emotion of the source (‘happiness’ instead of ‘sadness’). The inadequate scoring is a result of the WordNet taxonomy which causes the metric to equally treat both pairs (‘sad’, ‘happy’) and (‘cheerful’, ‘happy’) as non-synonym synsets and hence unmatched chunks.

The second problem with METEOR which may affect its efficacy in evaluating sentiment transfer relates to its weighting schema for function and non-function words. The following example clarifies the gravity of this problem:

### Example 3

Scores: [METEOR: 0.92]

- Hypothesis: “If he had blown himself up in your country, God would forgive him”
- Reference: “If he had blown himself up in your country, God would **not** forgive”.

The hypothesis in Example 3 is the translation output of Twitter’s built-in MT system for an Arabic tweet commenting on a terrorist attack<sup>1</sup>. The MT failure to translate the negation marker flips the sentiment of the author from ‘anger’ against the terrorist to ‘sympathy’. Despite this, the METEOR score is 0.92 which is within the highest upper bound, ranking it as a good translation. On the other hand, the METEOR score for a correct translation of sentiment with a negation marker is 0.93. The main culprit for this inaccurate scoring is the lexical weighting which causes the metric not to penalise the missing of a negation marker which produces a sentiment-critical error. Due to the grave consequences of such mistranslations, it becomes critical to have a sentiment-sensitive metric that is capable of spotting similar errors.

## 2.2 Embedding-based Metrics

**BERTScore** Recently embedding-based metrics have proven to achieve the highest performance in recent WMT shared tasks for quality metrics (e.g. Sellam et al. (2020); Lo (2020); Mukherjee et al. (2020)). We take BERTScore as a representative metric for this approach (Zhang et al., 2019). BERTScore computes a score based on a pair wise cosine similarity between the BERT contextual embeddings of the individual tokens for the hypothesis and the reference (Devlin et al., 2018). Accordingly, a BERTScore close to 1 indicates proximity in vector space and hence a good translation. The main problem with embedding-based metrics, such as BERTScore, is that antonyms contain similar distributional information since they usually occur in similar contexts. Example 4 illustrates this point:

### Example 4

Scores: [BERTScore: 0.85]

- Hypothesis: “What is this amount of anger; I don’t understand!”

<sup>1</sup><https://twitter.com/gaston810/status/673950532340465664>

- *Reference: “What is this amount of happiness, I don’t understand!”*

Example 4 shows the mistranslation produced by Twitter’s *Translate Tweet* tab of an Arabic tweet. Although the sentiment polarity is flipped in the candidate translation above, the hypothesis receives a BERTScore of 0.85 which indicates a high cosine similarity to the reference in vector space and hence a good translation. Clearly, the metric score is not comparable to a human perception of the emotion reflected by the source.

Figure 1 illustrates the reason behind this misleadingly high score. Figure 1 is a 2-D visualisation of BERT’s contextual embedding vectors for the hypothesis translation (in blue) and the reference (in red). Both sentences are very close in the embedding space due to the exact match of their individual tokens. The only mismatch is between the antonyms ‘happiness’ and ‘anger’. As shown in the figure, the pre-trained embedding vectors of the opposite polarity nouns are also quite close because of their common contextual information. An embedding metric such as BERTScore, therefore, may not penalise antonyms which typically occur in similar contexts.



Figure 1: Visualisation of Contextual Embeddings of Sentences in Example 4

Recently, there have been different approaches to overcome some distributional problems of contextual embeddings. Reimers and Gurevych (2019), for example, introduce SBert, a modification of the pretrained BERT network, which should mitigate the antonymy problem. They use Siamese network structures where the embeddings of similar sentence pairs are independently learned via two parallel transformer architectures. We measured how far this technique could solve the opposite sentiment

problem by measuring the SBert sentence similarity of the hypothesis and reference in Example 4. The cosine similarity of the SBert sentence embedding vectors for the hypothesis and reference in Example 4 reached 0.61. A correct translation of the reference, however, such as ‘What is all this cheerfulness, I don’t understand’ has a cosine similarity score of 0.79. This small similarity difference (0.18) would be misleading if taken as an evaluation of how far the sentiment poles in the two hypotheses are different. A more sentiment-targeted measure is needed for assessing mistranslations due to this antonymy problem.

In the following section, we conduct an experiment to quantify the divergence of the above mentioned metrics from the human perception of a proper sentiment transfer in a translated text.

### 3 Evaluation of Mistranslated Sentiment

We have shown in the previous section that three canonical MT evaluation metrics do not give a penalty proportional to sentiment-critical errors on segment-level by an MT online tool. In order to quantify how far the aforementioned quality metrics diverge or correlate with human judgement of sentiment transfer, we measure the performance of each metric on a dataset of tweets that had sentiment-critical translation errors. To compile this data, first we used the Twitter built-in translation system (Google API) to translate a dataset of tweets annotated for sentiment<sup>2</sup>. The source dataset amounted to  $\approx 7,000$  tweets in three languages: English, Arabic and Spanish. We translated the Arabic and Spanish into English and the English was translated into Spanish and Arabic.

To extract instances where the MT system failed to translate the sentiment correctly, we built an English sentiment-detection classifier by fine-tuning a Roberta XML model (Liu et al., 2019) on an English dataset of 23,000 tweets annotated for sentiment<sup>3</sup>. The English classifier was used to predict the sentiment of the Google API output for the translation of the Arabic and Spanish tweets into English and the English back-translation of the English tweets translated into Spanish and Arabic. The classifier’s predicted sentiment was compared to the gold standard emotion of the source text, and

<sup>2</sup>The dataset is collected from SemEval 2018 shared sentiment detection Task (Mohammad et al., 2018).

<sup>3</sup>The English dataset is collected from different sentiment detection tasks (Lo, 2019; Mohammad and Bravo-Marquez, 2017).

instances of discrepancy were extracted as potential mistranslations of sentiment. Finally, from the extracted instances, we manually built a translation quality evaluation dataset.

The quality evaluation dataset, henceforth QE, consisted of target tweets where the error is exclusively a mistranslation of the sentiment-carrying lexicon. In these tweets, the mistranslations either completely flip the sentiment polarity of the source tweet, similar to Examples 3 and 4 above, or transfer the same polarity but with a mitigated sentiment tone. The tweets with exclusive sentiment translation errors amounted to 300 tweets. We also added 100 tweet/translation pairs where the MT system transfers the correct sentiment. Reference translations of the QE dataset were created by native speakers of the respective source languages. Essentially, the reference translations aimed at correcting chunks that caused a distortion of the affect message and retained as many of the hypothesis n-grams as possible to detect how far each metric is sensitive to sentiment mismatching and not to the mismatch in other non-sentiment carrying words. The translators were also asked to assign a score to each pair of source-hypothesis tweet, where 1 is the poorest sentiment transfer and 10 is best sentiment transfer. The average scores of annotators were taken as the final human score<sup>4</sup>.

To quantify the ability of the three metrics explained in Section 2 to assess the transfer of sentiment in the QE dataset, we compared their scores of the translation hypotheses with the human judgement scores for sentiment transfer on the segment-level. We followed the WMT standard methods for evaluating quality metrics and used absolute Pearson correlation coefficient  $r$  and the Kendall correlation coefficient  $|\tau|$  to evaluate each metric’s performance against the human judgement. Figures 2 and 3 show heatmaps visualising the Pearson and Kendall correlation coefficients for the studied metrics and the human scores, respectively.

As seen from Figures 2 and 3, both BERTScore and METEOR achieve a better correlation with the human judgement than BLEU<sup>5</sup> which achieves only 0.16 and 0.12 Pearson and Kendall correlations, respectively. However, the relatively overall low correlations (max  $r = 0.39$  and max  $|\tau| = 0.27$ ) raise important doubts as to the reliability of these

<sup>4</sup>Annotators were computational linguists working on MT research.

<sup>5</sup>We use the Sacrebleu implementation of the BLEU score for all the experiments (Post, 2018).




Figure 2: Absolute Pearson correlations with segment-level human judgements for the QE dataset




Figure 3: Kendall correlations with segment-level human judgements for the QE dataset

accepted metrics for ranking MT systems which translate sentiments. Bearing in mind that 75% of the segments in the QE dataset have critical translation errors that seriously distort the sentiment, the low correlation results highlight the need for a sentiment-targeted measure that can improve a metric’s efficacy in capturing mistranslated sentiment by an MT system in real-life scenarios.

#### 4 Sentiment-Aware Measure (SAM) for Machine Translated UGC

In this section, we propose a new measure for assessing MT performance that takes into account the sentiment similarity between the MT system translation and the reference. This sentiment measure should be used as a fine-tuning tool to adjust a quality score in cases where it is used to assess the translation quality of sentiment-oriented text. The SAM score is calculated by using the SentiWord dictionary of prior polarities (Gatti et al., 2015). SentiWord is a sentiment lexicon that combines the high precision of manual lexica and the high coverage of automatic ones (covering 155,000 words). It is based on assigning a ‘prior polarity’ score for each lemma-POS in both SentiWordNet and a number of human-annotated sentiment lexica (Baccianella et al., 2010; Warriner et al., 2013). The

prior polarity is the out-of-context positive or negative score which a lemma-POS evokes. It is reached via an ensemble learning framework that combines several formulae where each lemma-POS is given the score that receives the highest number of votes from the different formulae. SentiWord prior polarity scores have been proven to achieve the highest correlation with human scores in sentiment analysis regression and classification SemEval tasks (Gatti et al., 2015).

We assume that our sentiment adjustment factor, SAM, is proportional to the distance between the sentiment scores of the unmatched words in the system translation and the reference, the higher the distance the greater the SAM adjustment. To calculate the SAM score, we designate the number of remaining mismatched words in the system translation and reference translation by  $m$  and  $n$ , respectively. We calculate the total SentiWord sentiment score for the lemma-POS of the mismatched words in the translation and reference sentences using a weighted average of the sentiment score of each mismatched lemma-POS. The weight of a hypothesis mismatched word  $w_h$  and a reference mismatched word  $w_r$  is calculated based on the sentiment score of its lemma-POS,  $s$ , as follows:

$$w_h^i = |s_i| \quad i = 1, 2, \dots, m. \quad (1)$$

$$w_r^i = |s_i| \quad i = 1, 2, \dots, n. \quad (2)$$

Then the total sentiment score for hypothesis  $S_h$  and reference  $S_r$  is given by:

$$S_h = \sum_{i=1}^m \alpha_i s_i, \quad \alpha_i = \frac{w_h^i}{\sum_{i=1}^m w_h^i} \quad (3)$$

$$S_r = \sum_{i=1}^n \beta_i s_i, \quad \beta_i = \frac{w_r^i}{\sum_{i=1}^n w_r^i} \quad (4)$$

The normalised SAM adjustment is given by:

$$p = \frac{|S_r - S_h|}{2} \quad (5)$$

and the translation final score will be given by:

$$\text{Score} = C_{hr} (1 - p) \quad (6)$$

where  $C_{hr}$  can be any metric’s matching score between a translation hypothesis and a reference segment. For this experiment, we will measure  $C_{hr}$  as the BLEU, METEOR and BERTScore scores. To illustrate how SAM score adjusts a metric score with respect to the transfer of sentiment, table 1

shows the SAM score adjustment for Examples 3 and 4 above.

As can be seen from Table 1, the metric score  $C_{hr}$  (here METEOR and BERTScore, respectively) is significantly reduced due to the sentiment distance between the mismatched words ( $w_h, w_r$ ) as well as their sentiment weights ( $S_h, S_r$ ). The reduced scores are more representative of the distortion of sentiment produced by the MT system in Examples 3 and 4. The SAM adjustment, therefore, is a targeted-measure that can fine-tune a metric according to the ‘sentiment-closeness’ of the translation and the reference<sup>6</sup>.

## 5 Assessing SAM performance

To check how far the SAM measure can improve the evaluation of sentiment, we assessed the performance of the three chosen metrics on the QE dataset utilised in the experiment in Section 3 with the SAM adjustment. Figures 4 and 5 show heatmaps of the Pearson and Kendall correlations of the SAM adjusted metrics with the human judgement in the QE dataset.




Figure 4: Absolute Pearson correlations with SAM Adjustment for the QE dataset

Compared to metric scores without the SAM adjustment, Figures 4 and 5 show that the combination of SAM with the three metrics consistently leads to an improvement in each metric’s correlation with the human judgement of a correct sentiment transfer in our dataset. Overall, BERTScore records the highest Pearson correlation coefficients and Kendall rank dependence (0.59 and 0.44, respectively). This means that it is better able to penalise critical translation errors in our QE dataset. Moreover, compared to their scores without SAM,

<sup>6</sup>More examples of SAM adjustments are in Appendix A

Table 1: Calculating the SAM adjustment for Examples 3 and 4

Example	$w_h$	$w_r$	$S_h$	$S_r$	$p$	$C_{hr}$	Score+SAM
3	him#a	not#r	0	-1.0	0.5	0.92	0.46
4	anger#n	happiness#n	-0.669	0.856	0.762	0.85	0.20




Figure 5: Kendall correlations with SAM adjustment for the QE dataset

both BERTScore and METEOR record 20% and 18% higher correlation, respectively. It is also worthy of notice, that although the Pearson correlation of BLEU improves from 0.16 to 0.33 with the SAM adjustment, the correlation is still relatively small. Knowing that BLEU score is usually the gate-keeping tool for accepting improvement in MT research, the results cast doubt on the efficacy of this non-semantic method to penalise sentiment-critical translation errors.

## 6 Error Analysis and Limitations

Human sentiment can be expressed in intricately subtle ways so that the mistranslation of the affect message is not necessarily reflected in divergence of polarity scores. We conducted an error analysis on instances where the SAM adjustment scores were not able to capture the MT’s failure to transfer the correct sentiment due to different linguistic phenomena. The first phenomenon we identified is related to structure shifting. For example, the sentiment distance between the source tweet ‘*I was saddened by him*’ and its mistranslation ‘*I made him sad*’ is very small despite the flipping of sentiment direction in the translation. The three metrics with and without the SAM adjustment failed to penalise this type of distortion in the QE dataset. Second, some words in the lexicon did not have a score representative of their sentiment weight. For example, most prepositions in the SentiWord lex-

icon are neutral, yet by checking the data, it was found out that a mistranslation of a preposition can distort the affect message. For example, the source tweet ‘*What is the benefit of me in this world*’ was mistranslated by the MT system as ‘*What is the benefit for me in this world*’; the wrong preposition causes the translation to lose the sad tone in the source tweet. Again, similar instances were not adequately measured with the three metrics with and without the SAM adjustment. Third, some nuanced sentiment-carrying words specific for the informal style of tweets caused a mistranslation of sentiment which was not captured by our approach. For instance, a one-word tweet referring to a political figure as ‘*Prick*’ was mistranslated as ‘*Sting*’. The source is a slang word used to refer to a mean, contemptible man. The translation wrongly received a high score since both ‘prick’ and ‘sting’ had very similar negative sentiment weights. The translation, however, fails to reflect the aggressive sentiment in the source tweet. The BLEU metric, on the other hand, succeeded in giving a penalty to similar short mistranslated tweets without the SAM adjustment.

One other limitation to the current approach for assessing the translation of sentiment is that it relies on an English sentiment lexicon. The applicability of this approach to other languages depends on the availability of a similar high-precision and high-coverage sentiment lexicon. We have overcome this problem by using the English backtranslations of the Arabic and Spanish tweets in the QE dataset. There are, however, multilingual translations of English sentiment lexica that are commonly used in the sentiment-analysis of non-English text. It remains to be tested how far a translated sentiment lexicon is capable of measuring sentiment transfer among different language pairs by the SAM scoring approach.

## 7 Conclusion

The most frequent scenario in which an MT system is used to translate sentiment-oriented text is for the translation of online UGC such as tweets, reviews or social media posts. The users of the MT online tool take these translations at face-value as there is

no human intervention for accuracy checking. It is important, therefore, to ensure the reliability of the MT system to accurately transfer the author’s affect message before it is used as an online tool. The reliability of an MT system with such big data is commonly measured by automatic metrics. We have shown that conventionally accepted metrics may not always be an optimum solution for assessing the translation of sentiment-oriented UGC. We presented an empirical approach to quantify the notion of sentiment transfer by an MT system and, more concretely, to modify automatic metrics such that its MT ranking comes closer to a human judgement of a poor or good translation of sentiment. Despite limitations to our approach, the SAM adjustment serves as a proxy to the complicated task of manually evaluating the translation of sentiment across different languages.

## Acknowledgements

This research work has been supported in part by the TranSent project at the Centre for Translation Studies (CTS) of the University of Surrey.

## References

- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In *Lrec*, volume 10, pages 2200–2204.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 conference on machine translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 376–380.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Lorenzo Gatti, Marco Guerini, and Marco Turchi. 2015. Sentiwords: Deriving a high precision and high coverage lexicon for sentiment analysis. *IEEE Transactions on Affective Computing*, 7(4):409–421.
- Yinuo Guo and Junfeng Hu. 2019. METEOR++ 2.0: Adopt Syntactic Level Paraphrase Knowledge into Machine Translation Evaluation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 501–506.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Chi-kiu Lo. 2019. YiSi-A unified semantic MT quality evaluation and estimation metric for languages with different levels of available resources. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 507–513.
- Chi-kiu Lo. 2020. Extended study on using pretrained language models and YiSi-1 for machine translation evaluation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 895–902.
- Qingsong Ma, Ondřej Bojar, and Yvette Graham. 2018. Results of the WMT18 metrics shared task: Both characters and embeddings achieve good performance. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 671–688.
- Nitika Mathur, Tim Baldwin, and Trevor Cohn. 2020a. Tangled up in BLEU: Reevaluating the Evaluation of Automatic Machine Translation Evaluation Metrics. *arXiv preprint arXiv:2006.06264*.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020b. Results of the WMT20 Metrics Shared Task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 688–725.
- Saif M. Mohammad and Felipe Bravo-Marquez. 2017. WASSA-2017 shared task on emotion intensity. In *Proceedings of the Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (WASSA)*, Copenhagen, Denmark.
- Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 Task 1: Affect in tweets. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, New Orleans, LA, USA.
- Ananva Mukherjee, Hema Ala, Manish Shrivastava, and Dipti Misra Sharma. 2020. Mee: An automatic metric for evaluation using embeddings for machine translation. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 292–299. IEEE.



- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing 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.
- Ted Pedersen, Siddharth Patwardhan, Jason Michelizzi, et al. 2004. Wordnet: Similarity-measuring the relatedness of concepts. In *AAAI*, volume 4, pages 25–29.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. *arXiv preprint arXiv:1804.08771*.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Ehud Reiter. 2018. A Structured Review of the Validity of BLEU. *Computational Linguistics*, 44(3):393–401.
- Hadeel Saadany and Constantin Orasan. 2020. [Is it great or terrible? preserving sentiment in neural machine translation of Arabic reviews](#). In *Proceedings of the Fifth Arabic Natural Language Processing Workshop*, pages 24–37, Barcelona, Spain (Online). Association for Computational Linguistics.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur P. Parikh. 2020. [Learning to evaluate translation beyond english: Bleurt submissions to the wmt metrics 2020 shared task](#).
- Kosuke Takahashi, Katsuhito Sudoh, and Satoshi Nakamura. 2020a. [Automatic machine translation evaluation using source language inputs and cross-lingual language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3553–3558, Online. Association for Computational Linguistics.
- Yujin Takahashi, Satoru Katsumata, and Mamoru Komachi. 2020b. [Grammatical error correction using pseudo learner corpus considering learner’s error tendency](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 27–32, Online. Association for Computational Linguistics.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior research methods*, 45(4):1191–1207.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. BERTScore: Evaluating text generation with Bert. *arXiv preprint arXiv:1904.09675*.

## Appendix A Examples for Mistranslated Tweets Measured with and without SAM adjustment

MT output	Reference	BLEU/METEOR/BERTScore	BLEU/METEOR/BERTScore+SAM
If he had blew himself up in your country, God would forgive him.	If he had blew himself up in your country, may God <u>not</u> forgive. <sup>1</sup>	0.70/0.83/0.96	0.35/0.41/0.48
Currently, I am <u>excited</u> about the one who criticizes those who are not happy at Eid.	Currently, I am <u>angry</u> about the one who criticizes those who are not happy at Eid.	0.77/0.52/0.89	0.23/0.15/0.26
My study skills are <u>awesome</u> .	My study skills are <u>terrible</u> .	0.40/0.41/0.85	0.11/0.11/0.27
May God bless you, the destroyer of Bahrain.	May God <u>not</u> bless you, the destroyer of Bahrain.	0.63/0.53/0.60	0.35/0.29/0.33
<u>My brother</u> , this match looks like the super match.	<u>I'm afraid</u> this match will look like the super match.	0.25/0.63/0.90	0.05/0.18/0.22

<sup>1</sup> Translation errors are underlined.