

BLEU Meets COMET: Combining Lexical and Neural Metrics Towards Robust Machine Translation Evaluation

Taisiya Glushkova^{1,3} Chrysoula Zerva^{1,3} André F. T. Martins^{1,2,3}

¹Instituto de Telecomunicações ²Unbabel

³Instituto Superior Técnico & LUMILIS (Lisbon ELLIS Unit)

{taisiya.glushkova, chrysoula.zerva, andre.t.martins}@tecnico.ulisboa.pt

Abstract

Although neural-based machine translation evaluation metrics, such as COMET or BLEURT, have achieved strong correlations with human judgements, they are sometimes unreliable in detecting certain phenomena that can be considered as critical errors, such as deviations in entities and numbers. In contrast, traditional evaluation metrics such as BLEU or CHRF, which measure lexical or character overlap between translation hypotheses and human references, have lower correlations with human judgements but are sensitive to such deviations. In this paper, we investigate several ways of combining the two approaches in order to increase robustness of state-of-the-art evaluation methods to translations with critical errors. We show that by using additional information during training, such as sentence-level features and word-level tags, the trained metrics improve their capability to penalize translations with specific troublesome phenomena, which leads to gains in correlation with human judgements and on recent challenge sets on several language pairs.¹

1 Introduction

Trainable machine translation (MT) evaluation models, such as COMET (Rei et al., 2020) and BLEURT (Sellam et al., 2020), generally achieve higher correlations with human judgments, thanks

to leveraging pretrained language models. However, they often fail at detecting certain types of errors and deviations from the source, for example related to translations of numbers and entities (Amrhein and Sennrich, 2022). As a result, their quality predictions are sometimes hard to interpret and not always trustworthy. In contrast, traditional lexical-based metrics, such as BLEU (Papineni et al., 2002) or CHRF (Popović, 2015)—despite their many limitations—are considerably more sensitive to these errors, due to their nature, and are also more interpretable, since the scores can be traced back to the character or n -gram overlap.

This paper investigates and compares methods that combine the strengths of neural-based and lexical approaches, both at the sentence level and at the word level. This is motivated by the findings of previous works, which demonstrate in detail that the COMET MT evaluation metric struggles to handle errors like deviation in numbers, wrong named entities in generated translations, deletions that exclude important content from the source sentence, insertions of extra words that are not present in the source sentences, and a few others (Amrhein and Sennrich, 2022; Alves et al., 2022). While data augmentation techniques alleviate the problem to some extent (Alves et al., 2022), the gains seem to be relatively modest. In this paper we investigate alternative methods that take advantage of lexical information and go beyond the use of various augmentation techniques and synthetic data.

We focus on increasing robustness of MT evaluation systems to certain types of critical errors. We experiment with the reference-based COMET metric, which has access to reference translations when producing quality scores. In order to make evaluation metrics more robust towards this type of errors, we consider and compare three different ways of

© 2023 The authors. This article is licensed under a Creative Commons 4.0 licence, no derivative works, attribution, CC-BY-ND.

¹Our code and data are available at: https://github.com/deep-spin/robust_MT_evaluation

incorporating information from lexical-based evaluation metrics into the neural-based COMET metric:

- Simply ensembling the sentence-level metrics;
- Using lexical-based sentence-level scores as additional features through a bottleneck layer in the COMET model;
- Enhancing the word embeddings computed by COMET for the generated hypothesis with word-level tags. We generate these word-level tags using the Levenshtein (sub)word alignment between the hypothesis and the reference tokens.

We compare these three strategies with the recent approach of (Alves et al., 2022), which generates synthetic data with injected errors from a large language model, and retrains COMET on training data that has been augmented with these examples. We assess both the correlation with human judgments and using the recently proposed DEMETR benchmark (Karpinska et al., 2022).

2 Related Work

Recently several challenge sets have been introduced, either within a scope of the WMT Metrics shared task (Freitag et al., 2022) or in general as a step towards implementing more reliable MT evaluation metrics: SMAUG (Alves et al., 2022) explores sentence-level multilingual data augmentation; ACES (Amrhein et al., 2022) is a translation accuracy challenge set that covers high number of different phenomena and language pairs, including a considerable number of low-resource ones; DEMETR (Karpinska et al., 2022) and HWTSC (Chen et al., 2022) aim at examining metrics ability to handle synonyms and to discern critical errors in translations; DFKI (Avramidis and Macketanz, 2022) employs a linguistically motivated challenge set for two language directions (German \leftrightarrow English).

Apart from purely focusing on improving robustness with augmentation of different phenomena, there are works that combine usage of synthetic data with other different methods. These methods use more fine-grained information—aiming at identifying both the position and the type of translation errors on given source-hypothesis sentence pairs (Bao et al., 2023). As another source of useful information, word-level supervision can be considered, which has proven to be beneficial in tasks of

quality estimation and MT evaluation (Rei et al., 2022a; Rei et al., 2022b).

There have been other attempts to add linguistic features to automatic MT evaluation metrics, e.g. incorporating information from a multilingual knowledge graph into BERTScore. (Wu et al., 2022) proposed a metric that linearly combines the results of BERTScore and bilingual named entity matching for reference-free machine translation evaluation. (Abdi et al., 2019) use an extensive set of linguistic features at word- and sentence- level to aid sentiment classification. Additionally, glass-box features extracted from the MT model have been used successfully in the quality estimation task (Fomicheva et al., 2020; Zerva et al., 2021; Wang et al., 2021). For the incorporation of different types of information to neural models early and late fusion is commonly used with benefits on multiple tasks and domains (Gadzicki et al., 2020; Fu et al., 2015; Baltrušaitis et al., 2018). To the best of our knowledge there have not been any attempts to combine the representations of neural metrics with external features obtained by lexical-based metrics.

Moreover, there are similar concerns regarding robustness of evaluation models in non-MT related tasks (Chen and Eger, 2022). In general, it is depicted that evaluation metrics perform rather well on standard evaluation benchmarks but are vulnerable and unstable to adversarial examples. The approaches investigated in our paper aim to address these limitations.

3 Combination of Neural and Lexical Metrics

In this section we describe the methods we investigated in order to infuse the COMET with information on lexical alignments between the MT hypothesis and the reference.

3.1 Metric ensembling

A simple way to combine neural and lexical-based metrics is through an ensembling strategy. To this end, we use a weighted ensemble of normalized BLEU, CHRf and COMET scores. The weights for each metric are tuned on the same development set used for training the COMET models discussed in this work (MQM WMT 2021) and presented in Appendix A. For normalisation we compute the mean and standard deviation to standardize the development set for each metric and we use the same mean and standard deviation values to standardize the

test-set scores.

3.2 Sentence-level lexical features

A simple ensemble is limited because it does not let the neural-based model *learn* the best way of including the information coming from the lexical metrics—for example, the degree of additional information brought by the lexical metrics might depend on the particular input.

Therefore, we experiment with a more sophisticated approach, where the lexical scores are incorporated in the COMET architecture as additional features that are mapped to each instance in the data, allowing the system to learn how to best take advantage of these features. To this end, we adopt a late fusion approach, employing a bottleneck layer to combine the lexical and neural features. The use of a bottleneck layer for late fusion in deep neural architectures has been used successfully across tasks, especially for multimodal fusion or fusion of features with vast differences in dimensionality (Petridis et al., 2017; Guo et al., 2018; Ding et al., 2022). In our implementation, the bottleneck layer is inserted between two feed-forward layers in the original COMET architecture (see Fig. 1), implemented in a similar manner as in (Moura et al., 2020; Zerva et al., 2022) (see App. A).

3.3 Word-level lexical features

While the sentence-level features allow the model to account for lexical overlap, there is still no word-level information. Instead, we propose to leverage the inferred alignments between the MT hypothesis and the reference words. To that end we adopt the Translation Edit Rate (TER) alignment procedure that calculates the edit distance (cost) between the translation and reference sentence. This alignment, produced with the Levenshtein dynamic programming algorithm, identifies the minimal subset of MT words that would need to be changed (modified, inserted, or deleted) in order to reach the correct translation (reference) (Snover et al., 2009b). TER-based alignments have been widely used to evaluate translations with respect to post-edits (HTER) in automated post-editing as well as other generation tasks (Snover et al., 2009a; Elliott and Keller, 2014; Gupta et al., 2018; Bhattacharyya et al., 2022). Recently, providing word-level supervision using binary quality tags inferred via Multidimensional Quality Metrics (MQM) error annotations, proved to be beneficial for MT evaluation (Rei et al., 2022a).

In this work, for simplicity, we opted for calculating the alignments not on a word but on a sub-word level, employing the same tokenization convention used by the COMET encoder.² This allows to associate a quality OK / BAD tag to each sub-word unit of the MT hypothesis input vector.

We then incorporate these quality tags to the original input for each translation sample which consists of a triplet $\langle s, t, r \rangle$, where s is a source text, t is a machine translated text, and r is a reference translation. To leverage the estimated quality tags in the COMET architecture, we encode the tags as a sequence of special tokens, w , and learn separate embeddings for the OK / BAD tokens. We can thus encode the quality tag sequence and obtain a word quality vector \vec{w} and then compute the sum $\vec{\sigma} = \vec{t} \oplus \vec{w}$ for the sequence. We then extend the pooling layer of COMET by adding both the \vec{w} and $\vec{\sigma}$ representations (see the architecture in Fig. 2).

4 Experimental Design

The main focus of our experiments is to investigate how the robustness of the MT evaluation models can be improved and how the proposed settings compare to each other and to a data augmentation approach proposed by (Alves et al., 2022). Our comparisons address the correlation with human judgments and recent robustness benchmarks on MT evaluation datasets (§4).

We follow (Amrhein and Sennrich, 2022) – we use COMET (v1.0) (Rei et al., 2020) as the underlying architecture for our MT evaluation models and focus on making it more robust.

Human Judgements Data We consider two types of human judgments: direct assessments (DA) (Graham et al., 2013) and multi-dimensional quality metric scores (MQM) (Lommel et al., 2014). For training, we use WMT 2017–2020 data from the Metrics shared task (Freitag et al., 2021b) with direct assessment (DA) annotations (see App. C). For development and test, we use the MQM annotations of the WMT 2021 and 2022 datasets, respectively³.

Challenge Sets Data Furthermore, we evaluate our models using two challenge sets: DEMETR (Karpinska et al., 2022) and ACES (Amrhein et al., 2022).

²We specifically used the `XLRobertaTokenizerFast` Huggingface implementation with truncation and default `max_length`.

³We opted for DA annotations to train due to the limited availability of MQM data

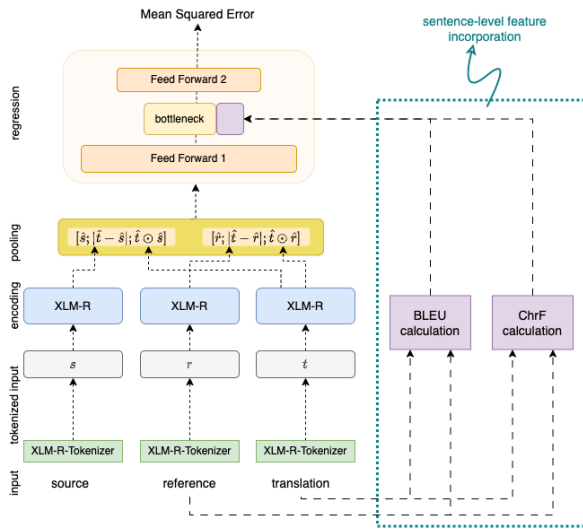


Figure 1: The architecture of the COMET model with incorporated sentence-level lexical features.

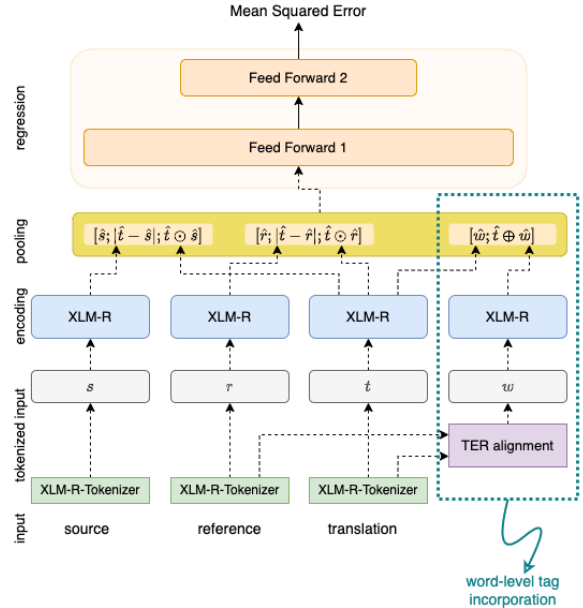


Figure 2: The architecture of the COMET model with incorporated word-level lexical features.

- DEMETR is a diagnostic dataset with 31K English examples (translated from 10 source languages) created for evaluating the sensitivity of MT evaluation metrics to 35 different linguistic perturbations spanning semantic, syntactic, and morphological error categories. Each example in DEMETR consists of (1) a sentence in one of 10 source languages, (2) an English translation written by a human translator, (3) a machine translation produced by Google Translate, and (4) an automatically perturbed version of the Google Translate output which introduces exactly one mistake (semantic, syntactic, or typographical).
- ACES is a translation accuracy challenge set based on the MQM ontology. It consists of 36,476 examples covering 146 language pairs and representing 68 phenomena. This challenge set consists of synthetically generated adversarial examples, examples from repurposed contrastive MT test sets, and manually annotated examples.

Both of these challenge sets allow measuring the sensitivity of the proposed approaches to various phenomena and assess their overall robustness.

Augmentation We compare our methods against the multilingual data augmentation approach SMAUG⁴ proposed by (Alves et al., 2022). Specif-

⁴The code is available at <https://github.com/Unbabel/smaug>.

ically, we use transformations that focus on deviations in named entities and numbers since these are identified as the major weaknesses of COMET (Amrhein and Sennrich, 2022).

Models In the experiments that follow, we use as baseline the vanilla COMET architecture trained on WMT2017–2020 (COMET). We compare this baseline against the model trained with augmented data and our proposed approaches:

- **COMET + aug:** COMET model trained on a mixture of original and augmented WMT2017–2020 data, where the percentage of the augmented data is 40%. We use the code provided by the authors of SMAUG and apply their choice of hyperparameters, including the optimal percentage of the augmented data.
- **Ensemble:** The weighted mean of BLEU, CHRf and COMET normalized scores, where the weights are optimized on the development set (MQM 2021) with regards to the Kendall’s tau correlations.
- **COMET + SL-feat.:** The combination of COMET and scores obtained from other metrics, BLEU and CHRf, that are used as sentence-level (SL) features in a late fusion manner.
- **COMET + WL-tags:** The combination of COMET and word-level OK / BAD tags that correspond to the subwords of the translation

hypothesis.

Evaluation For evaluation and analysis we:

1. Compute standard correlation metrics on segment-level between predicted scores and human judgements: Pearson r , Spearman ρ and Kendall’s tau;
2. Use challenge sets, specifically DEMETR and ACES, to analyse the robustness of MT Evaluation systems to critical errors and specific perturbations.

For the challenge sets, we measure the ability of the evaluation metric to rank the correct translations higher than the incorrect ones by computing the official Kendall’s tau-like correlation as proposed in previous WMT Metrics shared tasks (Freitag et al., 2022; Ma et al., 2019):

$$\tau = \frac{\text{Concordant} - \text{Discordant}}{\text{Concordant} + \text{Discordant}}, \quad (1)$$

where the “Concordant” is the number of times a metric assigns a higher score to the “better” hypothesis and “Discordant” is the number of times a metric assigns a higher score to the “worse” hypothesis.

5 Results and Discussion

In this section, we show results for the aforementioned methods, specifically the correlations with MQM annotations from WMT 2022 Metrics shared task for 3 high-resource language pairs (English \rightarrow German, English \rightarrow Russian, Chinese \rightarrow English) in four domains: Conversation, E-commerce, News and Social. In addition, we discuss evaluation results obtained on two challenge sets.

5.1 Correlation with Human Judgements

Overall, by looking at Table 1 we can see that the more sophisticated techniques of using additional information, whether it is lexical-based scores used as features, word-level tags based on token alignments or synthetically augmented data, outperform the simple weighted average (ensemble) approach. These findings are further supported when calculating performance for the Pearson r and Spearman ρ coefficients, shown in Tables 9 and 10 respectively in the Appendix B.

Across all proposed methods, we observe that **COMET + aug** and **COMET + SL-feat.** have relatively similar performance. In contrast, adding

word-level tags (**COMET + WL-tags**) based on alignments between the translation hypothesis and the reference seems to give a considerable gain in results compared to the baseline **COMET** and the other approaches.

Another interesting observation is that the improvement in correlations can be especially noticed in ZH-EN language pair across all domains for **COMET + WL-tags** model. Overall, we found that adding the word-level quality supervision provides the most consistent benefits in performance. However, since our main motivation is to address robustness to specific errors, the correlations with MQM annotations serve primarily as a confirmation of the potential of the proposed methods; we provide a more detailed performance analysis over the multiple error types of different challenge sets in the next section.

5.2 Results on Challenge Sets

5.2.1 DEMETR

For DEMETR we analyse results on two levels of granularity: (1) performance over the full challenge set, which is calculated via Kendall’s tau and presented in Table 2 which shows Kendall’s tau-like correlations per language pair; and (2) performance depending on error severity, which is presented in and Table 3 and shows accuracy on detecting different types of DEMETR perturbations for lexical and neural-based metrics, bucketed by error severity (baseline, critical, major, and minor errors).

We can observe that both the sentence- and word-level features outperform data augmentation methods, with the word-level method being the best on average and for the majority of language pairs. These findings indicate that the subword quality tags enable the model to attend more to the perturbations of the high quality data, hence better distinguishing the bad from the good translations of the same source.

One of the key findings from Table 3 is that the model which uses word-level information consistently outperforms the other methods across almost all severity buckets, with the exception of “critical” error bucket. In combination with the findings on the ACES challenge set (see section 5.2.2), it seems that investigating approaches which target more nuanced and complex error phenomena that lead to

⁵For the statistical significance over correlations r we use Williams’ test and Fisher $r - to - z'$ transform: $f(r) = \frac{1}{2} \ln \frac{1+r}{1-r}$ to calculate significance over the macro-averages, with $p \leq 0.01$.

		BLEU	CHRF	COMET	ENSEMBLE	COMET+aug	COMET+SL-feat.	COMET+WL-tags
EN-DE	Conversation	0.201	0.257	0.308	0.309	0.296	0.310	0.314
	E-commerce	0.179	0.212	0.326	0.318	0.311	0.322	0.322
	News	0.167	0.202	0.361	0.356	0.330	0.355	0.369
	Social	0.130	0.168	0.297	0.292	0.277	0.294	0.293
EN-RU	Conversation	0.140	0.175	0.305	0.304	0.328	0.298	0.328
	E-commerce	0.202	0.221	0.372	0.371	0.382	0.369	0.391
	News	0.125	0.164	0.373	0.367	0.366	0.384	0.370
	Social	0.152	0.132	0.305	0.304	0.330	0.332	0.349
ZH-EN	Conversation	0.125	0.160	0.283	0.282	0.295	0.283	0.298
	E-commerce	0.174	0.187	0.326	0.325	0.342	0.335	0.357
	News	0.046	0.042	0.270	0.261	0.291	0.276	0.292
	Social	0.162	0.190	0.319	0.316	0.313	0.315	0.330
AVG		0.150	0.176	0.321	0.317	0.322	0.323	0.334 [†]

Table 1: Kendall’s tau correlation on high resource language pairs using the MQM annotations for Conversation, E-commerce, News and Social domains collected for the WMT 2022 Metrics Task. **Bold** numbers indicate the best result for each domain in each language pair. † in the averaged scores indicates statistically significant difference to the other metrics ⁵.

	BLEU	CHRF	COMET	ENSEMBLE	COMET+aug	COMET+SL-feat.	COMET+WL-tags
ZH-EN	0.505	0.684	0.818	0.855	0.817	0.866	0.872
DE-EN	0.655	0.802	0.909	0.926	0.917	0.942	0.957
HI-EN	0.616	0.768	0.900	0.92	0.925	0.929	0.945
JA-EN	0.521	0.722	0.850	0.883	0.83	0.907	0.891
PS-EN	0.533	0.703	0.818	0.88	0.775	0.863	0.877
RU-EN	0.552	0.724	0.898	0.91	0.894	0.950	0.949
CZ-EN	0.541	0.755	0.875	0.917	0.863	0.87	0.920
FR-EN	0.664	0.794	0.892	0.915	0.926	0.945	0.951
ES-EN	0.516	0.704	0.877	0.899	0.877	0.91	0.934
IT-EN	0.601	0.774	0.912	0.924	0.906	0.936	0.945
AVG	0.57	0.743	0.875	0.903	0.873	0.912	0.924 [†]

Table 2: Kendall’s tau-like correlation per language pair on DEMETR challenge set. **Bold** values indicate the best performance per language pair. † in the averaged scores indicates statistically significant difference to the other metrics.

Metric	Base	Crit.	Maj.	Min.	All
<i>lexical-based metrics</i>					
BLEU	100.0	79.33	83.76	72.6	78.52
CHRF	100.0	90.79	90.85	80.83	87.16
<i>neural-based metrics</i>					
ENSEMBLE	100.0	96.87	92.91	93.77	95.14
COMET	99.3	95.77	91.04	92.18	93.74
+ aug	98.6	95.54	91.66	92.06	93.65
+ SL-feat.	99.3	96.95	93.56	94.64	95.59
+ WL-tags	99.2	96.48	93.9	96.36	96.2

Table 3: Accuracy on DEMETR perturbations for both lexical-based and neural-based metrics, shown bucketed by error severity (base, critical, major, and minor errors), including a micro-average across all perturbations.

critical errors could further improve performance of neural metrics.

5.2.2 ACES

To analyse general, high-level, performance trends of the lexical and proposed approaches on the ACES challenge set, we report Kendall’s tau correlation and the “ACES - Score” as proposed by

(Amrhein et al., 2022), which is a weighted combination of the 10 top-level categories in the ACES ontology:

$$\text{ACES-Score} = \text{sum} \left\{ \begin{array}{l} 5 * \tau_{\text{addition}} \\ 5 * \tau_{\text{omission}} \\ 5 * \tau_{\text{mistranslation}} \\ 5 * \tau_{\text{overtranslation}} \\ 5 * \tau_{\text{undertranslation}} \\ 1 * \tau_{\text{untranslated}} \\ 1 * \tau_{\text{do not translate}} \\ 1 * \tau_{\text{real-world knowledge}} \\ 1 * \tau_{\text{wrong language}} \\ 0.1 * \tau_{\text{punctuation}} \end{array} \right\} \quad (2)$$

The weights in this formula correspond to the recommended values in the MQM framework (Freitag et al., 2021a): *weight* = 5 for major, *weight* = 1 for minor and *weight* = 0.1 for fluency/punctuation errors. The ACES-Score results can be seen in the last row of Table 4.

	BLEU	CHRF	COMET	ENSEMBLE	COMET+aug	COMET+SL-feat.	COMET+WL-tags
<i>major (weight = 5)</i>							
addition	0.748	0.644	0.349	0.367	0.52	0.443	0.427
omission	0.427	0.784	0.704	0.828	0.706	0.724	0.666
mistranslation	-0.296	0.027	0.186	0.216	0.255	0.148	0.189
overtranslation	-0.838	-0.696	0.27	0.176	0.308	0.086	0.304
undertranslation	-0.856	-0.592	0.08	-0.044	0.2	-0.18	0.12
<i>minor (weight = 1)</i>							
untranslated	0.786	0.928	0.709	0.894	0.58	0.618	0.686
do not translate	0.58	0.96	0.88	0.9	0.78	0.9	0.84
real-world knowl.	-0.906	-0.307	0.195	0.176	0.202	0.109	0.162
wrong language	0.659	0.693	0.071	0.052	0.159	0.185	0.087
<i>fluency/punctuation (weight = 0.1)</i>							
punctuation	0.658	0.803	0.328	0.699	0.377	0.323	0.339
ACES-Score	-2.89	3.189	9.833	9.807	11.704	7.949	10.339

Table 4: Kendall’s tau-like correlations for 10 top-level categories in ACES challenge set.

	BLEU	CHRF	COMET	ENSEMBLE	COMET+aug	COMET+SL-feat.	COMET+WL-tags
EN-XX	0.034	0.329	0.201	0.340	0.256	0.183	0.206
XX-EN	-0.37	-0.046	0.283	0.26	0.329	0.222	0.285
XX-YY	-0.124	0.097	0.105	0.115	0.204	0.088	0.104
AVG	-0.153	0.127	0.196	0.238	0.263 [†]	0.164	0.198

Table 5: Kendall’s tau-like correlation on ACES challenge set. † in the averaged scores indicates statistically significant difference to the other metrics.

Overall, as the ACES challenge set contains a larger set of translation errors, and goes beyond simple perturbations to more nuanced error categories such as real-world knowledge and discourse-level errors, we can see that the performance scores and best metrics vary largely depending on the category. Interestingly, CHRF seems to outperform other metrics especially in the categories that do not relate so much to replacements in the reference translation, but rather relate to fully or partially wrong language (or punctuation) use. We note that these seem to be largely cases that are not frequently found in MT training data, nor are they considered in previously proposed data augmentation approaches, which could explain why neural metrics are outperformed by baseline surface-level metrics, even under investigated robustness modifications. Hence, there seems to be room for further improvements in incorporating surface-based information in neural metrics and enabling them to pay more attention to n-gram overlap. Instead, for the error categories that depend on other perturbations, we can see that all robustness oriented modifications to COMET improve the performance compared to the vanilla model, with augmentation achieving significantly higher Kendall’s tau correlations.

When looking at the overall picture and focusing

on the ACES-Score which weights the errors by the severity of the errors there seem to be only two methods that outperform the baseline COMET model, namely COMET + aug and COMET + WL-tags, which achieve the best and second best ACES-Score respectively. Since these two approaches are orthogonal to each other, it seems that a promising direction for future work is to explore options for combining the two methods.

Note that the overall behavior of lexical and neural-based metrics corroborates the findings presented in the original paper. We can confirm that in our experiments the worst performing metric is also BLEU, which is expected. However, it is hard to highlight the best performing metric based only on the ACES-Score, the purpose of this analysis is more so to find any interesting trends or any particular issues that some methods are handling better than the others.

Since the ACES dataset encompasses a high number of LPs, we aggregate the results into three groups, EN-XX (out-of-English), XX-EN (into-English) and XX-YY (LPs without English). We also report the balanced average across all language pairs (AVG). Results in Table 5 show that methods which include augmented data during training achieve higher performance compared to

other proposed options. As for additional sentence-level or word-level information, **COMET + WL-tags** slightly improves performance of the baseline COMET across EN-XX and XX-EN aggregations and beats the approach that uses SL-features.

6 Conclusion and Future Work

In this paper, we presented several approaches that use interpretable string-based metrics to improve the robustness of recent neural-based metrics such as COMET. There are various ways of combining these methods together: ensembling metrics, incorporating sentence-level features, or using word-level information coming from alignments between the hypothesis and the reference. We observed that adding small changes to the architecture of COMET, either by using sentence-level features based on BLEU and CHRf scores, or by incorporating word-level tags for the hypothesis, can lead to competitive performance gains. To showcase the effectiveness of our proposed approaches, we evaluated them on the most recent MQM test set that covers multiple domains and language pairs, as well as on the challenge sets that were introduced in the WMT 2022 Metrics shared task, with encouraging results.

It is likely that our proposed approaches are complementary to each other, as well as to the data augmentation method we are comparing against (COMET+aug). An interesting direction for future work is to study further the impact of using word-level tags of the hypothesis in other ways not covered in this paper, e.g., in combination with augmentation approaches.

Acknowledgements

This work was supported by the European Research Council (ERC StG DeepSPIN 758969), by EU's Horizon Europe Research and Innovation Actions (UTTER, contract 101070631), by P2020 project MAIA (LISBOA-01-0247- FEDER045909), by the Portuguese Recovery and Resilience Plan through project C645008882-00000055 (NextGenAI, Center for Responsible AI) and Fundação para a Ciência e Tecnologia through contract UIDB/50008/2020.

References

- [Abdi et al.2019] Abdi, Asad, Siti Mariyam Shamsuddin, Shafaatunnur Hasan, and Jalil Piran. 2019. Deep learning-based sentiment classification of evaluative
- [Alves et al.2022] Alves, Duarte, Ricardo Rei, Ana C Farinha, José G. C. de Souza, and André F. T. Martins. 2022. Robust MT evaluation with sentence-level multilingual augmentation. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 469–478, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Amrhein and Sennrich2022] Amrhein, Chantal and Rico Sennrich. 2022. Identifying Weaknesses in Machine Translation Metrics Through Minimum Bayes Risk Decoding: A Case Study for COMET. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*, Online, November. Association for Computational Linguistics.
- [Amrhein et al.2022] Amrhein, Chantal, Nikita Moghe, and Liane Guillou. 2022. ACES: Translation accuracy challenge sets for evaluating machine translation metrics. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 479–513, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Avramidis and Macketanz2022] Avramidis, Eleftherios and Vivien Macketanz. 2022. Linguistically motivated evaluation of machine translation metrics based on a challenge set. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 514–529, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Baltrušaitis et al.2018] Baltrušaitis, Tadas, Chaitanya Ahuja, and Louis-Philippe Morency. 2018. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443.
- [Bao et al.2023] Bao, Keqin, Yu Wan, Dayiheng Liu, Baosong Yang, Wenqiang Lei, Xiangnan He, Derek F Wong, and Jun Xie. 2023. Towards fine-grained information: Identifying the type and location of translation errors. *arXiv preprint arXiv:2302.08975*.
- [Bhattacharyya et al.2022] Bhattacharyya, Pushpak, Rajen Chatterjee, Markus Freitag, Diptesh Kanojia, Matteo Negri, and Marco Turchi. 2022. Findings of the wmt 2022 shared task on automatic post-editing. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 109–117.
- [Chen and Eger2022] Chen, Yanran and Steffen Eger. 2022. Menli: Robust evaluation metrics from natural language inference. *arXiv preprint arXiv:2208.07316*.
- [Chen et al.2022] Chen, Xiaoyu, Daimeng Wei, Hengchao Shang, Zongyao Li, Zhanglin Wu,

- Zhengzhe Yu, Ting Zhu, Mengli Zhu, Ning Xie, Lizhi Lei, Shimin Tao, Hao Yang, and Ying Qin. 2022. Exploring robustness of machine translation metrics: A study of twenty-two automatic metrics in the WMT22 metric task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 530–540, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Ding et al.2022] Ding, Ning, Sheng-wei Tian, and Long Yu. 2022. A multimodal fusion method for sarcasm detection based on late fusion. *Multimedia Tools and Applications*, 81(6):8597–8616.
- [Elliott and Keller2014] Elliott, Desmond and Frank Keller. 2014. Comparing automatic evaluation measures for image description. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 452–457.
- [Fomicheva et al.2020] Fomicheva, Marina, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:539–555.
- [Freitag et al.2021a] Freitag, Markus, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021a. Experts, errors, and context: A large-scale study of human evaluation for machine translation. *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- [Freitag et al.2021b] Freitag, Markus, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021b. 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, November. Association for Computational Linguistics.
- [Freitag et al.2022] Freitag, Markus, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. Results of WMT22 metrics shared task: Stop using BLEU – neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Fu et al.2015] Fu, Zhikang, Bing Li, Jun Li, and Shuhua Wei. 2015. Fast film genres classification combining poster and synopsis. In *Intelligence Science and Big Data Engineering. Image and Video Data Engineering: 5th International Conference, IScIDE 2015, Suzhou, China, June 14-16, 2015, Revised Selected Papers, Part 1 5*, pages 72–81. Springer.
- [Gadzicki et al.2020] Gadzicki, Konrad, Razieh Khamsehashari, and Christoph Zetsche. 2020. Early vs late fusion in multimodal convolutional neural networks. In *2020 IEEE 23rd international conference on information fusion (FUSION)*, pages 1–6. IEEE.
- [Graham et al.2013] Graham, Yvette, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. Continuous measurement scales in human evaluation of machine translation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41, Sofia, Bulgaria, August. Association for Computational Linguistics.
- [Guo et al.2018] Guo, Lili, Longbiao Wang, Jianwu Dang, Linjuan Zhang, and Haotian Guan. 2018. A feature fusion method based on extreme learning machine for speech emotion recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2666–2670. IEEE.
- [Gupta et al.2018] Gupta, Ankush, Arvind Agarwal, Prawaan Singh, and Piyush Rai. 2018. A deep generative framework for paraphrase generation. In *Proceedings of the aaai conference on artificial intelligence*, volume 32.
- [Karpinska et al.2022] Karpinska, Marzena, Nishant Raj, Katherine Thai, Yixiao Song, Ankita Gupta, and Mohit Iyyer. 2022. DEMETR: Diagnosing evaluation metrics for translation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9540–9561, Abu Dhabi, United Arab Emirates, December. Association for Computational Linguistics.
- [Lommel et al.2014] Lommel, Arle, Hans Uszkoreit, and Aljoscha Burchardt. 2014. Multidimensional quality metrics (mqm): A framework for declaring and describing translation quality metrics. *Tradumàtica*, (12):0455–463.
- [Ma et al.2019] Ma, Qingsong, 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, August. Association for Computational Linguistics.
- [Moura et al.2020] Moura, João, Miguel Vera, Daan van Stigt, Fabio Kepler, and André F. T. Martins. 2020. IST-unbabel participation in the WMT20 quality estimation shared task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 1029–1036, Online, November. Association for Computational Linguistics.
- [Papineni et al.2002] Papineni, Kishore, 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, July. Association for Computational Linguistics.
- [Petridis et al.2017] Petridis, Stavros, Zuwei Li, and Maja Pantic. 2017. End-to-end visual speech recognition with lstms. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 2592–2596. IEEE.
- [Popović2015] Popović, Maja. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal, September. Association for Computational Linguistics.
- [Rei et al.2020] Rei, Ricardo, 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, November. Association for Computational Linguistics.
- [Rei et al.2022a] Rei, Ricardo, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022a. COMET-22: Unbabel-IST 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Rei et al.2022b] Rei, Ricardo, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022b. CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid), December. Association for Computational Linguistics.
- [Sellam et al.2020] Sellam, Thibault, 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, July. Association for Computational Linguistics.
- [Snover et al.2009a] Snover, Matthew, Nitin Madnani, Bonnie Dorr, and Richard Schwartz. 2009a. Fluency, adequacy, or hter? exploring different human judgments with a tunable mt metric. In *Proceedings of the fourth workshop on statistical machine translation*, pages 259–268.
- [Snover et al.2009b] Snover, Matthew G, Nitin Madnani, Bonnie Dorr, and Richard Schwartz. 2009b. Ter-plus: paraphrase, semantic, and alignment enhancements to translation edit rate. *Machine Translation*, 23:117–127.
- [Wang et al.2021] Wang, Ke, Yangbin Shi, Jiayi Wang, Yuqi Zhang, Yu Zhao, and Xiaolin Zheng. 2021. Beyond glass-box features: Uncertainty quantification enhanced quality estimation for neural machine translation. *arXiv preprint arXiv:2109.07141*.
- [Wu et al.2022] Wu, Zhanglin, Min Zhang, Ming Zhu, Yinglu Li, Ting Zhu, Hao Yang, Song Peng, and Ying Qin. 2022. Kg-bertscore: Incorporating knowledge graph into bertscore for reference-free machine translation evaluation. In *11th International Joint Conference on Knowledge Graphs, IJCKG2022. To be published*.
- [Zerva et al.2021] Zerva, Chrysoula, Daan Van Stigt, Ricardo Rei, Ana C Farinha, Pedro Ramos, José GC de Souza, Taisiya Glushkova, Miguel Vera, Fabio Kepler, and André FT Martins. 2021. Ist-unbabel 2021 submission for the quality estimation shared task. In *Proceedings of the Sixth Conference on Machine Translation*, pages 961–972.
- [Zerva et al.2022] Zerva, Chrysoula, Taisiya Glushkova, Ricardo Rei, and André F. T. Martins. 2022. Disentangling uncertainty in machine translation evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8622–8641, Abu Dhabi, United Arab Emirates, December. Association for Computational Linguistics.

A Model Implementation and Parameters

Table 8 shows the hyperparameters used to train the following prediction models: **COMET**, **COMET + SL-feat.** and **COMET + WL-tags**. For the baseline we used the code available at <https://github.com/Unbabel/COMET> and we trained the model on WMT17-WMT20 DA data (in the table we refer to it as **COMET**).

For the ENSEMBLE we tune three weights on the development set with grid search, by optimizing Kendall tau correlations (see Table 6).

	BLEU	CHRF	COMET
weights	0.02513	0.04523	0.92965

Table 6: Tuned weights on the MQM 2021 set for the weighted ensemble.

The bottleneck size parameter for **COMET + SL-feat.** model was tuned using development set. This set covers three language pairs (English \rightarrow German, English \rightarrow Russian, Chinese \rightarrow English) and two domains (ted and newstest). Kendall tau correlation was computed over the whole dataset without considering different domains separately (see Table 7).

	64	128	256	512
EN-DE	0.223	0.216	0.217	0.225
EN-RU	0.305	0.279	0.275	0.281
ZH-EN	0.319	0.330	0.325	0.315
AVG	0.282	0.275	0.272	0.274

Table 7: Kendall’s tau-like correlation per language pair on development set for different bottleneck sizes. **Bold** values indicate the best performance per language pair.

B Correlation with Human Judgements

We present here results on MQM 2022 set for Pearson and spearman correlations (see Tables 9 and 10 accordingly). We can see that especially for Spearman ρ the findings are aligned with the findings on Kendall tau correlations. Instead, for the Pearson r which is more sensitive to outliers, we can see that the augmentation method outperforms the feature-based modifications.

C Training Data Statistics

The combined WMT training data (from 2017 to 2020) has 950069 segments and covers the following language pairs (total number is 32): Cs-En, De-Cs, De-En, De-Fr, En-Cs, En-De, En-Et, En-Fi, En-Gu, En-Ja, En-Kk, En-Lt, En-Lv, En-Pl, En-Ru,

En-Ta, En-Tr, En-Zh, Et-En, Fi-En, Fr-De, Gu-En, Ja-En, Kk-En, Km-En, Lt-En, Pl-En, Ps-En, Ru-En, Ta-En, Tr-En, Zh-En.

Hyperparameter	COMET	COMET + SL-feat.	COMET + WL-tags
Encoder Model	XLM-R (large)	XLM-R (large)	XLM-R (large)
Optimizer	Adam	Adam	Adam
No. frozen epochs	0.3	0.3	0.3
Learning rate	3e-05	3e-05	3e-05
Encoder Learning Rate	1e-05	1e-05	1e-05
Layerwise Decay	0.95	0.95	0.95
Batch size	4	4	4
Loss function	Mean squared error	Mean squared error	Mean squared error
Dropout	0.15	0.15	0.15
Hidden sizes	[3072, 1024]	[3072, 1024]	[3072, 1024]
Encoder Embedding layer	Frozen	Frozen	Frozen
Bottleneck layer size	-	64	-
FP precision	32	32	32
No. Epochs (training)	2	2	2

Table 8: Hyperparameters used to train different prediction methods.

		BLEU	CHRf	COMET	ENSEMBLE	COMET + aug	+ SL-feat.	+ WL-tags
EN-DE	Conversation	0.228	0.285	0.371	0.376	0.378	0.379	0.400
	Ecommerce	0.173	0.222	0.376	0.373	0.380	0.383	0.341
	News	0.220	0.260	0.521	0.521	0.492	0.506	0.526
	Social	0.172	0.220	0.367	0.367	0.375	0.382	0.351
EN-RU	Conversation	0.155	0.185	0.372	0.369	0.418	0.350	0.400
	Ecommerce	0.249	0.287	0.488	0.488	0.510	0.507	0.481
	News	0.169	0.230	0.469	0.467	0.464	0.477	0.448
	Social	0.213	0.143	0.324	0.328	0.371	0.343	0.385
ZH-EN	Conversation	0.160	0.206	0.340	0.338	0.370	0.343	0.358
	Ecommerce	0.220	0.230	0.391	0.391	0.438	0.400	0.440
	News	0.097	0.078	0.340	0.334	0.383	0.364	0.359
	Social	0.161	0.177	0.351	0.347	0.358	0.343	0.373
AVG		0.185	0.210	0.393	0.392	0.411	0.398	0.405

Table 9: Pearson correlation on high resource language pairs using the MQM annotations for Conversation, Ecommerce, News and Social domains collected for the WMT 2022 Metrics Task. **Bold** numbers indicate the best result for each domain in each language pair.

		BLEU	CHRf	COMET	ENSEMBLE	COMET + aug	+ SL-feat.	+ WL-tags
EN-DE	Conversation	0.262	0.337	0.401	0.403	0.385	0.404	0.409
	Ecommerce	0.235	0.278	0.421	0.411	0.403	0.416	0.417
	News	0.224	0.273	0.478	0.472	0.438	0.471	0.486
	Social	0.173	0.222	0.389	0.383	0.361	0.386	0.384
EN-RU	Conversation	0.183	0.230	0.400	0.397	0.427	0.389	0.428
	Ecommerce	0.276	0.303	0.502	0.501	0.514	0.499	0.528
	News	0.171	0.224	0.499	0.492	0.490	0.514	0.495
	Social	0.212	0.186	0.425	0.423	0.455	0.460	0.483
ZH-EN	Conversation	0.166	0.211	0.375	0.369	0.385	0.370	0.389
	Ecommerce	0.241	0.259	0.449	0.448	0.467	0.459	0.487
	News	0.063	0.057	0.364	0.352	0.393	0.373	0.394
	Social	0.219	0.256	0.424	0.421	0.418	0.419	0.439
AVG		0.202	0.236	0.427	0.423	0.428	0.430	0.445

Table 10: Spearman correlation on high resource language pairs using the MQM annotations for Conversation, Ecommerce, News and Social domains collected for the WMT 2022 Metrics Task. **Bold** numbers indicate the best result for each domain in each language pair.