

MT Metrics Correlate with Human Ratings of Simultaneous Speech Translation

Dominik Macháček¹ and Ondřej Bojar¹ and Raj Dabre²

Charles University, Faculty of Mathematics and Physics,
Institute of Formal and Applied Linguistics¹

National Institute of Information and Communications Technology, Kyoto, Japan²

¹{machacek,bojar}@ufal.mff.cuni.cz, ²raj.dabre@nict.go.jp

Abstract

There have been several meta-evaluation studies on the correlation between human ratings and offline machine translation (MT) evaluation metrics such as BLEU, chrF2, BERTSCORE and COMET. These metrics have been used to evaluate simultaneous speech translation (SST) but their correlations with human ratings of SST, which has been recently collected as Continuous Ratings (CR), are unclear. In this paper, we leverage the evaluations of candidate systems submitted to the English-German SST task at IWSLT 2022 and conduct an extensive correlation analysis of CR and the aforementioned metrics. Our study reveals that the offline metrics are well correlated with CR and can be reliably used for evaluating machine translation in simultaneous mode, with some limitations on the test set size. We conclude that given the current quality levels of SST, these metrics can be used as proxies for CR, alleviating the need for large scale human evaluation. Additionally, we observe that correlations of the metrics with translation as a reference is significantly higher than with simultaneous interpreting, and thus we recommend the former for reliable evaluation.

1 Introduction

The current approach to evaluate simultaneous speech translation (SST, [Cho and Esipova, 2016](#); [Ma et al., 2019](#)) systems that have text as the output modality is to use automatic metrics which are designed for offline text-to-text machine translation (MT), alongside other measures for latency and stability. Researchers tend to use offline metrics, such as BLEU ([Papineni et al., 2002](#)), chrF2 ([Popović, 2017](#)), BERTSCORE ([Zhang et al., 2020](#)), COMET ([Rei et al., 2020](#)) and others ([Freitag et al., 2022](#)) in SST despite no explicit evidence that they correlate with human ratings.

However, simultaneous speech-to-text translation has different characteristics compared to offline text-to-text MT. For example, when the users

are following subtitles in real-time, they have limited time for reading and comprehension as they cannot fully control the reading pace by themselves. Therefore, they may be less sensitive to subtle grammar and factual flaws than while reading a text document without any time constraints. The human evaluation of SST should therefore reflect the simultaneity. The users may also prefer brevity and simplicity over verbatim word-for-word translation. Even if the reference is brief and simpler than the original, there may be lots of variants that the BLEU score and other MT metrics may not evaluate as correct.

Furthermore, SST and MT differ in their input modalities. MT sources are assumed to originate as texts, while the SST source is a speech given in a certain situation, accompanied by para-linguistic means and specific context knowledge shared by the speaker and listener. Transcribing speech to text for use in offline evaluation of SST may be limiting.

In this paper, we aim to determine the suitability of automatic metrics for evaluating SST. To this end, we analyze the results of the simultaneous speech translation task from English to German at IWSLT 2022 ([Anastasopoulos et al., 2022](#)), where we calculate the correlations between MT metrics and human judgements in simultaneous mode. There are five competing systems and human interpreting that are manually rated by bilingual judges in a simulated real-time event. Our studies show that BLEU does indeed correlate with human judgements of simultaneous translations under the same conditions as in offline text-to-text MT: on a sufficiently large number of sentences. Furthermore, chrF2, BERTSCORE and COMET exhibit similar but significantly larger correlations. To the best of our knowledge, we are the first to explicitly establish the correlation between automatic offline metrics with human SST ratings, indicating that they may be safely used in SST evaluation in

the currently achieved translation quality levels.

Additionally, we statistically compare the metrics with translation versus interpreting reference, and we recommend the most correlating one: translation reference and COMET metric, with BERTSCORE and chrF2 as fallback options.

We publish the code for analysis and visualisations that we created in this study.¹ It is available for further analysis and future work.

2 Related Work

We replicate the approach from text-to-text MT research (e.g. Papineni et al., 2002) that examined the correlation of MT metrics with human judgments. The strong correlation is used as the basis for taking the metrics as reliable. As far as we know, we are the first who apply this approach to SST evaluation in simultaneous mode.

In this paper, we analyze four metrics that represent the currently used or recommended (Freitag et al., 2022) types of MT metrics. BLEU and chrF2 are based on lexical overlap and are available for any language. BERTSCORE (Zhang et al., 2020) is based on embedding similarity of a pre-trained BERT language model. COMET (Rei et al., 2020) is a neural metric trained to estimate the style of human evaluation called Direct Assessment (Graham et al., 2015). COMET requires sentence-to-sentence aligned source, translation and reference in the form of texts, which may be unavailable in some SST use-cases; then, other metric types may be useful. Another fact is that BERTSCORE and COMET are available only for a limited set of languages.

3 Human Ratings in SST

As far as we know, the only publicly available collection of simultaneous (not offline) human evaluation of SST originates from IWSLT 2022 (Salesky et al., 2022) English-to-German Simultaneous Translation Task, which is described in “Findings” (Anastasopoulos et al., 2022, see highlights of it we discuss in Appendix A). The task focused on speech-to-text translation and was reduced to translation of individual sentences. The segmentation of the source audio to sentences was provided by organizers, and not by the systems themselves. The source sentence segmentation that was used in human evaluation was gold (oracle). It only approximates a realistic setup where the

segmentation would be provided by an automatic system, e.g. Tsiamas et al. (2022), and may be partially incorrect and cause more translation errors compared to the gold segmentation.

The simultaneous mode in Simultaneous Translation Task means that the source is provided gradually, one audio chunk at a time. After receiving each chunk, the system decides to either wait for more source context, or produce target tokens. Once the target tokens are generated, they can not be rewritten.

The participating systems are submitted and studied in three latency regimes: low, medium and high. It means that the maximum Average Lagging (Ma et al., 2019) between the source and target on validation set must be 1, 2 or 4 seconds in a “computationally unaware” simulation where the time spent by computation, and not by waiting for context, is not counted. One system in low latency did not pass the latency constraints (see Findings, page 44, numbered 141), but it is manually evaluated regardless.

Computationally unaware latency was one of the main criteria in IWSLT 2022. It means that the participants did not need to focus on a low latency implementation, as it is more of a technical and hardware issue than a research task. However, the subtitle timing in manual evaluation was created in a way such that waiting for the first target token was dropped, and then it continued with computationally aware latency.

3.1 Continuous Rating (CR)

Continuous Rating (CR, Javorský et al., 2022; Macháček and Bojar, 2020) is a method for human assessment of SST quality in a simulated online event. An evaluator with knowledge of the source and target languages watches a video (or listens to an audio) document with subtitles created by the SST system which is being evaluated. The evaluator is asked to continuously rate the quality of the translation by pressing buttons with values 1 (the worst) to 4 (the best). Each evaluator can see every document only once, to ensure one-pass access to the documents, as in a realistic setup.

CR is analogous to Direct Assessment (Graham et al., 2015), which is a method of human text-to-text MT evaluation in which a bilingual evaluator expresses the MT quality by a number on a scale. It is natural that individual evaluators have different opinions, and thus it is a common practice to

¹github.com/ufal/MT-metrics-in-SimST

have multiple evaluators evaluate the same outputs and then report the mean and standard deviation of evaluation scores, or the results of statistical significance tests that compare the pairs of candidate systems and show how confident the results are.

Javorský et al. (2022) showed that CR relates well to comprehension of foreign language documents by SST users. Using CR alleviates the need to evaluate comprehension by factual questionnaires that are difficult to prepare, collect and evaluate. Furthermore, Javorský et al. (2022) show that bilingual evaluators are reliable.

Criteria of CR In IWSLT 2022, the evaluators were instructed that the primary criterion in CR should be meaning preservation (or adequacy), and other aspects such as fluency should be secondary. The instructions do not mention readability due to output segmentation frequency or verbalizing non-linguistic sounds such as “laughter”, despite the system candidates differ in these aspects.

3.2 Candidate Systems

Automatic SST systems There are 5 evaluated SST systems: FBK (Gaido et al., 2022), NAIST (Fukuda et al., 2022), UPV (Iranzo-Sánchez et al., 2022), HW-TSC (Wang et al., 2022), and CUNI-KIT (Polák et al., 2022).

Human Interpreting In order to compare the state-of-the-art SST with human reference, the organizers hired one expert human interpreter to simultaneously interpret all the test documents. Then, they employed annotators to transcribe the voice into texts. The annotators worked in offline mode. The transcripts were then formed as subtitles including the original interpreter’s timing and were used in CR evaluation the same way as SST. However, human interpreters use their own segmentation to translation units so that they often do not translate one source sentence as one target sentence. There is no gold alignment of the translation sentences to interpreting chunks. The alignment has to be resolved before applying metrics to interpreting.

3.3 Evaluation Data

There are two subsets of evaluation data used in IWSLT22 En-De Simultaneous Translation task. The “Common” subset consists of TED talks of the native speakers. See the description in Findings on page 9 (numbered as 106). The “Non-Native” subset consists of mock business presentations of European high school students (Macháček et al.,

2019), and of presentations by representatives of European supreme audit institutions. This subset is described in Findings on page 39 (numbered page 136). The duration statistics of audio documents in both test sets are in Findings in Table 17 on page 48 (numbered 145).

4 Correlation of CR and MT Metrics

In this section, we study the correlation of CR and MT metrics BLEU, chrF2, BERTSCORE and COMET. We measure it on the level of documents, and not on the test set level, increasing the number of observations for significance tests. There are 60 evaluated documents (17 in the Common subset and 43 in Non-Native) and 15 system candidates (5 systems, each in 3 latency regimes), which yields 900 data points.

We discovered that CUNI-KIT system outputs are tokenized, while the others are detokenized. Therefore, we first detokenized CUNI-KIT outputs. Then, we removed the final end of sequence token (</s>) from the outputs of all systems. Finally, we calculated BLEU and chrF2 using sacreBLEU (Post, 2018), BERTSCORE and COMET. See Appendix B for metric details and signatures.

In total, there are 1584 rating sessions of 900 candidate document translations. Each candidate document translation is rated either twice with different evaluators, once, or not at all. We aggregate the individual rating clicks in each rating session by plain average (CR definition in Appendix C) to get the CR scores. Then, we average the CR of the same documents and candidate translations, and we correlate it with MT metrics.




Figure 1: Averaged document CR vs COMET on both Common and Non-Native subsets.

Averaged document ratings					
subsets	num.	BLEU	chrF2	BERTS.	COMET
both	823	0.65	0.73	0.77	0.80
Common	228	<i>0.42</i>	0.63	0.68	0.76
Non-Native	595	0.70	0.70	0.73	0.75

All document ratings					
subsets	num.	BLEU	chrF2	BERTS.	COMET
both	1584	0.61	0.68	0.71	0.73
Common	441	<i>0.37</i>	<i>0.57</i>	0.60	0.68
Non-Native	1143	0.64	0.64	0.66	0.67

Table 1: Pearson correlation coefficients for CR vs MT metrics BLEU, chrF2, BERTSCORE and COMET for averaged document ratings by all 5 SST systems and 3 latency regimes (upper), and all ratings (lower). When the coefficient is less than 0.6 (in *italics*), the correlation is not considered as strong. Significance values are $p < 0.01$ in all cases, meaning strong confidence.

Correlation Results In Table 1, we report correlation coefficients with and without averaging, together with the number of observations. Figure 1 displays the relation between CR and COMET.

Pearson correlation is considered as strong if the coefficient is larger than 0.6 (Evans, 1996). The results show strong correlation (above 0.65) of CR with BLEU, chrF2, BERTSCORE and COMET at the document level on both test subsets. When we consider only one subset, the correlation is lower, but still strong for chrF2, BERTSCORE and COMET (0.63, 0.68 and 0.76, resp.). It is because the Common subset is generally translated better than Non-Native, so with only one subset, the points span a smaller part of the axes and contain a larger proportion of outliers.

The strong correlation is not the case of BLEU on the Common subset where the Pearson coefficient is 0.42. We assume it is because BLEU is designed for use on a larger test set, but we use it on short single documents. However, BLEU correlates with chrF2 and COMET (0.81 and 0.62 on the Common subset). BLEU also correlates with CR on the level of test sets, as reported in Findings in the caption of Table 18 (page 48, numbered 145).

We conclude that with the current overall levels of speech translation quality, BLEU, chrF2, BERTSCORE and COMET can be used for reliable assessment of human judgement of SST quality at least on the level of test sets. chrF2, BERTSCORE and COMET are reliable also at the document level.

Translation vs Interpreting Reference There is an open question whether SST should rather mimic offline translation, or simultaneous interpreting. As

metric	reference	alignment	corr.
COMET	TRANSL	SENT	0.80
COMET	TRANSL	SINGLESEQ	0.79
COMET	TRANSL+INTP	SINGLESEQ	0.79
BERTSCORE	TRANSL	SENT	0.77
BERTSCORE	TRANSL+INTP	SENT+MWER	0.77
COMET	INTP	SINGLESEQ	0.77
BERTSCORE	TRANSL+INTP	SINGLESEQ	0.76
BERTSCORE	TRANSL	SINGLESEQ	0.75
chrF2	TRANSL+INTP	SENT+MWER	0.73
BLEU	TRANSL+INTP	SINGLESEQ	0.73
chrF2	TRANSL	SENT	0.73
chrF2	TRANSL+INTP	SINGLESEQ	0.72
chrF2	TRANSL	SINGLESEQ	0.72
BLEU	TRANSL	SINGLESEQ	0.71
COMET	INTP	MWER	0.71
BERTSCORE	INTP	SINGLESEQ	0.69
BLEU	TRANSL+INTP	SENT+MWER	0.68
chrF2	INTP	SINGLESEQ	0.66
BLEU	TRANSL	SENT	0.65
chrF2	INTP	MWER	0.65
BLEU	INTP	SINGLESEQ	0.65
BERTSCORE	INTP	MWER	0.60
BLEU	INTP	MWER	0.58

Table 2: Pearson correlation of metric variants to averaged CR on both subsets, ordered from the most to the least correlating ones. Lines indicate “clusters of significance”, i.e. boundaries between groups where all metric variants significantly differ from all in the other groups, with $p < 0.05$ for dashed line and $p < 0.1$ for dotted line. See the complete pair-wise comparison in Appendix D.

Macháček et al. (2021) discovered, translation may be more faithful, word-for-word, but also more complex to perceive by target audience. Simultaneous interpreting, on the other hand, tends to be brief and simpler than offline translation. However, it may be less fluent and less accurate. Therefore, we consider human translation (TRANSL) and transcript of simultaneous interpreting (INTP) as two possible references, and also test multi-reference metrics with both.

Since interpreting is not sentence-aligned to SST candidate translations, we consider two alignment methods: single sequence (SINGLESEQ), and mWERSegmenter (Matusov et al., 2005, MWER). SINGLESEQ method means that we concatenate all the sentences in the document to one single sequence, and then apply the metric on it, as if it was one sentence. mWERSegmenter is a tool for aligning translation candidates to reference, if their sentence segmentation differs. It finds the alignment with the minimum WER when comparing tokens in aligned segments. For translation, we also apply the default sentence alignment (SENT).

In Table 2, we report the correlations of metric,

reference and alignment variants and their significance, with more details in Appendix D.

4.1 Recommendations

Taking CR as the golden truth of human quality, we make the following recommendations of the most correlating metric, reference and sentence alignment method for SST evaluation.

Which metric? COMET, because it correlates significantly better with CR than BERTSCORE does. From the fall back options, chrF2 should be slightly preferred over BLEU.

Which reference? The metrics give significantly higher correlations with CR with translations than with interpreting as a reference. Difference between translation reference and two references (TRANSL+INTP) is insignificant. Therefore, we recommend translation as a reference for SST.

Which alignment method? With an unaligned reference, COMET and BERTSCORE correlate significantly more with SINGLESEQ than with MWER, probably because the neural metrics are trained on full, complete sentences, which are often split to multiple segments by mWERSegmenter. chrF2 correlates insignificantly better with MWER than with SINGLESEQ.

5 Conclusion

We found correlation of offline MT metrics to human judgements of simultaneous speech translation. The most correlating and thus preferred metric is COMET, followed by BERTSCORE and chrF2. We recommend text translation reference over interpreting, and single sequence alignment for neural, and mWERSegmenter for n -gram metrics.

6 Limitations

The data that we analyzed are limited to only one English-German language pair, 5 SST systems from IWSLT 2022, and three domains. All the systems were trained in the standard supervised fashion on parallel texts. They do not aim to mimic interpretation with shortening, summarization or redundancy reduction, and they do not use document context. The used MT metrics are good for evaluating individual sentence translations and that is an important, but not the only subtask of SST. We assume that some future systems created with a different approach may show divergence of CR and the offline MT metrics.

Furthermore, we used only one example of human interpreting. A precise in-depth study of human interpretations is needed to re-assess the recommendation of translation or interpreting as reference in SST.

Acknowledgements

We are thankful to Dávid Javorský and Peter Polák for their reviews.

This research was partially supported by the grants 19-26934X (NEUREM3) of the Czech Science Foundation, SVV project number 260 698, and 398120 of the Grant Agency of Charles University.

References

- Antonios Anastasopoulos, Loïc Barrault, Luisa Bentivogli, Marcelly Zanon Boito, Ondřej Bojar, Roldano Cattoni, Anna Currey, Georgiana Dinu, Kevin Duh, Maha Elbayad, Clara Emmanuel, Yannick Estève, Marcello Federico, Christian Federmann, Souhir Gahbiche, Hongyu Gong, Roman Grundkiewicz, Barry Haddow, Benjamin Hsu, Dávid Javorský, Věra Kloudová, Surafel Lakew, Xutai Ma, Prashant Mathur, Paul McNamee, Kenton Murray, Maria Nădejde, Satoshi Nakamura, Matteo Negri, Jan Niehues, Xing Niu, John Ortega, Juan Pino, Elizabeth Salesky, Jiatong Shi, Matthias Sperber, Sebastian Stüker, Katsuhito Sudoh, Marco Turchi, Yogesh Virkar, Alexander Waibel, Changhan Wang, and Shinji Watanabe. 2022. [Findings of the IWSLT 2022 evaluation campaign](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 98–157, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Kyunghyun Cho and Masha Esipova. 2016. [Can neural machine translation do simultaneous translation?](#) *CoRR*, abs/1606.02012.
- James D. Evans. 1996. *Straightforward Statistics for the Behavioral Sciences*. Pacific Grove: Brooks/Cole Pub.
- Markus Freitag, 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*, pages 46–68, Abu Dhabi. Association for Computational Linguistics.
- Ryo Fukuda, Yuka Ko, Yasumasa Kano, Kosuke Doi, Hirotaka Tokuyama, Sakriani Sakti, Katsuhito Sudoh, and Satoshi Nakamura. 2022. [NAIST simultaneous speech-to-text translation system for IWSLT 2022](#). In *Proceedings of the 19th International Conference on*

- Spoken Language Translation (IWSLT 2022)*, pages 286–292, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Marco Gaido, Sara Papi, Dennis Fucci, Giuseppe Fiameni, Matteo Negri, and Marco Turchi. 2022. [Efficient yet competitive speech translation: FBK@IWSLT2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 177–189, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Yvette Graham, Timothy Baldwin, and Nitika Mathur. 2015. [Accurate evaluation of segment-level machine translation metrics](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1183–1191, Denver, Colorado. Association for Computational Linguistics.
- Javier Iranzo-Sánchez, Javier Jorge Cano, Alejandro Pérez-González-de Martos, Adrián Giménez Pastor, Gonçal Garcés Díaz-Munío, Pau Baquero-Arnal, Joan Albert Silvestre-Cerdà, Jorge Civera Saiz, Albert Sanchis, and Alfons Juan. 2022. [MLLP-VRAIN UPV systems for the IWSLT 2022 simultaneous speech translation and speech-to-speech translation tasks](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 255–264, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Dávid Javorský, Dominik Macháček, and Ondřej Bojar. 2022. [Continuous rating as reliable human evaluation of simultaneous speech translation](#). In *Proceedings of the Seventh Conference on Machine Translation*, pages 154–164, Abu Dhabi. Association for Computational Linguistics.
- Mingbo Ma, Liang Huang, Hao Xiong, Renjie Zheng, Kaibo Liu, Baigong Zheng, Chuanqiang Zhang, Zhongjun He, Hairong Liu, Xing Li, Hua Wu, and Haifeng Wang. 2019. [STACL: Simultaneous translation with implicit anticipation and controllable latency using prefix-to-prefix framework](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3025–3036, Florence, Italy. Association for Computational Linguistics.
- Dominik Macháček and Ondřej Bojar. 2020. [Presenting simultaneous translation in limited space](#). In *Proceedings of the 20th Conference Information Technologies - Applications and Theory (ITAT 2020)*, Hotel Tyrapol, Oravská Lesná, Slovakia, September 18-22, 2020, volume 2718 of *CEUR Workshop Proceedings*, pages 34–39. CEUR-WS.org.
- Dominik Macháček, Jonáš Kratochvíl, Tereza Vojtěchová, and Ondřej Bojar. 2019. A speech test set of practice business presentations with additional relevant texts. In *Statistical Language and Speech Processing*, pages 151–161, Cham. Springer International Publishing.
- Dominik Macháček, Matúš Žilinec, and Ondřej Bojar. 2021. [Lost in Interpreting: Speech Translation from Source or Interpreter?](#) In *Proc. Interspeech 2021*, pages 2376–2380.
- Evgeny Matusov, Gregor Leusch, Oliver Bender, and Hermann Ney. 2005. [Evaluating machine translation output with automatic sentence segmentation](#). In *Proceedings of the Second International Workshop on Spoken Language Translation*, Pittsburgh, Pennsylvania, USA.
- 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.
- Peter Polák, Ngoc-Quan Pham, Tuan Nam Nguyen, Danni Liu, Carlos Mullov, Jan Niehues, Ondřej Bojar, and Alexander Waibel. 2022. [CUNI-KIT system for simultaneous speech translation task at IWSLT 2022](#). In *Proceedings of the 19th International Conference on Spoken Language Translation (IWSLT 2022)*, pages 277–285, Dublin, Ireland (in-person and online). Association for Computational Linguistics.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. 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, Brussels, Belgium. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [Unbabel’s participation in the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 911–920, Online. Association for Computational Linguistics.
- Elizabeth Salesky, Marcello Federico, and Marta Costajussà, editors. 2022. [Proceedings of the 19th International Conference on Spoken Language Translation \(IWSLT 2022\)](#). Association for Computational Linguistics, Dublin, Ireland (in-person and online).
- Ioannis Tsiamas, Gerard I. Gállego, José A. R. Fonollosa, and Marta R. Costa-jussà. 2022. [SHAS: Approaching optimal Segmentation for End-to-End Speech Translation](#). In *Proc. Interspeech 2022*, pages 106–110.
- Minghan Wang, Jiabin Guo, Yinglu Li, Xiaosong Qiao, Yuxia Wang, Zongyao Li, Chang Su, Yimeng Chen, Min Zhang, Shimin Tao, Hao Yang, and Ying Qin. 2022. [The HW-TSC’s simultaneous speech translation system for IWSLT 2022 evaluation](#). In *Proceedings of the 19th International Conference on Spoken*

Language Translation (IWSLT 2022), pages 247–254, Dublin, Ireland (in-person and online). Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

A Highlights of IWSLT22 Findings

The Findings of IWSLT22 (Anastasopoulos et al., 2022) are available in PDF. The most up-to-date version (version 2) is 61 pages long.² We highlight the relevant parts of Findings with page numbers in Table 3 so that we can refer to them easily.

Note that findings are a part of the conference proceedings (Salesky et al., 2022) as a chapter in a book. The order of findings pages in PDF does not match the page numbers at the footers.

Also note that in Section 2.4 on page 4 (in PDF, 101 in Proceedings), there is a description of MLLP-VRAIN which corresponds to the system denoted as UPV in all other tables and figures.

B Metric Signatures

BLEU and chrF2 SacreBLEU metric signature is case:mixedleff:noltok:13alsmooth:explversion:2.3.1.

For BERTSCORE, we used F1 with signature bert-base-multilingual-cased_L9_no-idf_version=0.3.12(hug_trans=4.23.1)_fast-tokenizer.

We use COMET model wmt20-comet-da (Rei et al., 2020). For multi-reference COMET, we run the model separately with each reference and average the scores.

The standard way of using mWERSegmenter is to segment candidate translation according to reference. However, COMET requires aligned source as one of the inputs, and mWERSegmenter can not align it because it is in other language. For COMET INTP MWER variant, we therefore aligned interpreting to translation, which is already aligned to source. For the other metrics with INTP MWER, we aligned translation candidate to interpreting, which is the standard way.

C Aggregating Continuous Ratings

We revisited the processing of the individual collected clicks on the rating buttons into the aggregate score of Continuous Rating.

²<https://aclanthology.org/2022.iwslt-1.10v2.pdf>

We found two definitions that can yield different results in certain situations: (1) The rating (as clicked by the evaluator) is valid at the instant time point when the evaluator clicked the rating button. The final score is the average of all clicks, each click has the equal weight. We denote this interpretation as CR .

(2) The rating is assigned to the time interval from the click time to the next click, or between the last click and the end of the document. The length of the interval is considered in averaging. The final score is the average of ratings weighted by interval lengths when the rating is valid. We denote this interpretation as CRi .³

To express them rigorously, let us have a document of duration T , and n ratings (r_i, t_i) , where $i \in \{1, \dots, n\}$ is an index, $r_i \in \{1, \dots, 4\}$ is the rated value and $0 \leq t_1 < \dots < t_n \leq T$ are times when the ratings were recorded.

Then, the definitions are as follows:

$$CR = \frac{1}{n} \sum_{i=1}^n r_i$$

$$CRi = \frac{1}{T - t_1} \left(\sum_{i=1}^{n-1} (t_{i+1} - t_i) r_i + (T - t_n) r_n \right)$$

If the judges press the rating buttons regularly, with a uniform frequency, then both definitions give equal scores. Otherwise, the CR and CRi may differ and may yield even opposite conclusions. For example, pressing “1” twelve times in one minute, then “4” and then waiting for one minute results in different scores: $CR = 1.2$, $CRi = 2$.


To examine the relationship between these definitions, we counted CR and CRi for each annotation of each document in the evaluation campaign. The results are in Figure 2 where we observe correlation between the two definitions. The Pearson correlation coefficient is 0.98, which indicates a very strong correlation.

Summary Based on the correlation score we observed, we conclude that both definitions are interchangeable, and any of them can be used in further analysis.

³Other interpretations are also conceivable, for instance assuming that the rating applies to a certain time before the click and then till the next judgement.

marker	PDF page	numbered page	description
Section 2	3-5	100-102	Simultaneous Speech Translation Task
Figure 1	6	103	Quality-latency trade-off curves
Section 2.6.1	5	102	Description of human evaluation
Figure 5	8	105	Manual scores vs BLEU (plot)
Two Test Sets (paragraph)	39	136	Non-Native subset
Test data (paragraph)	9	106	Common (native) subset of test data
Automatic Evaluation Results	44	141	Latency and BLEU results (table)
A1.1 (appendix)	38-39	135-136	Details on human evaluation
Table 17	48	145	Test subsets duration
Table 18	48	145	Manual scores and BLEU (table)

Table 3: Relevant parts of IWSLT22 Findings (<https://aclanthology.org/2022.iwslt-1.10v2.pdf>) for End-De Simultaneous Speech Translation task and human evaluation.



separately are lower and the differences along the diagonal are less significant. We explain it by the fact that in smaller data set, there is larger impact of noise.

Figure 2: Relation between weighted interval averaging of continuous rating (CRi, y-axis) and average of all ratings (CR, x-axis) for each annotation of each document (blue data points).

D Pairwise Metrics Comparison

We test the statistical difference of correlations with Steiger’s method.⁴ The method takes into account the number of data points and the fact that all three compared variables correlate, which is the case of the MT metrics that are applied on the same texts. We use two-tailed test.

We applied the test on all pairs of metric variants. The results for both subsets are in Figure 3. Figure 4 displays results on the Common subset, and Figure 5 for the Non-Native subset. These results are analogous to those in Table 1 in Section 4. The correlation scores for the two subsets treated

⁴<https://github.com/psinger/CorrelationStats/>

