

# An Analysis on Automated Metrics for Evaluating Japanese-English Chat Translation

Andre Rusli, Makoto Shishido  
Tokyo Denki University  
{20udc91@ms, shishido@mail}.dendai.ac.jp

## 概要

This paper analyses how traditional baseline metrics, such as BLEU and TER, and neural-based methods, such as BERTScore and COMET, score several NMT models' performance on chat translation and how these metrics perform when compared to human-annotated scores. The results show that for ranking NMT models in chat translations, all metrics seem consistent in deciding which model outperforms the others. This implies that traditional baseline metrics, which are faster and simpler to use, can still be helpful. On the other hand, when it comes to better correlation with human judgment, neural-based metrics outperform traditional metrics, with COMET achieving the highest correlation with the human-annotated score on a chat translation. However, we show that even the best metric struggles when scoring English translations from sentences with anaphoric zero-pronoun in Japanese.

## 1 Introduction

In recent years, machine translation has grown faster to provide better and more inclusive systems for many language directions. One of the most recent projects by the research team at Facebook tries to accommodate every language so that no language is left behind. Until a few years ago, most translation systems focused only on the sentence level, unable to understand the sentence's contexts which often exist in previous sentences. Many reports have demonstrated this to be inadequate. Several studies have shown that human translators outperform NMT models when the target language translation must consider the context at the document level [1, 2]. To improve this, the research community proposed various methods and improvements to build models that could use the source- and target-side contexts [3, 4] when performing

translations. By doing so, models could handle more challenging tasks which require context understanding, such as chat translations.

For Japanese NMT, various methods have also been proposed to tackle critical challenges, especially in translating Japanese conversational texts to English. One of the most prominent challenges in discourse is the anaphoric zero-pronoun, which arises from pronouns such as subjects, objects, and possessive cases omitted from conversational sentences. Anaphoric zero-pronoun is not exclusive to Japanese but is one of the most challenging languages to resolve [5, 6]. In contrast, while the number of works of research for training better and more capable NMT models keeps increasing rapidly, progress in machine translation (MT) evaluation has been struggling to keep up. Many works of research in the MT research community still rely on traditional metrics that are outdated but widely adopted as the standard.

Even though human judgment is still considered the best metric to measure translation quality, it is considerably more expensive and time-consuming. Therefore, automatic metrics are an indispensable part of machine translation evaluation. They provide immediate feedback during MT system development and serve as the primary metric to report the quality of MT systems. Accordingly, the reliability of metrics is critical to progress in MT research. Historically, metrics for evaluating the quality of MT have relied on assessing the similarity between an MT-generated hypothesis and a human-generated reference translation in the target language. Traditional metrics, such as BLEU [7] and TER [8], have focused on basic, lexical-level features, such as counting the number of matching n-grams between the MT hypothesis and the reference translation. On the other hand, more recently proposed metrics are primarily based on pre-trained Transformers-based language models, such as BERTScore [9] and COMET [10], to calculate the

similarity of translated sentences with their target. Additionally, similarity-based metrics built using neural-based methods can help evaluate human-generated texts, such as content scoring systems [11], often found in foreign language learning settings.

This paper compares how different evaluation metrics perform on chat translation from Japanese to English. Then we analyze and emphasize the importance of using suitable metrics for measuring the performance of our translation systems depending. Additionally, we provide datasets containing the translation results of the Japanese-English Business Scene Dialogue corpus by three NMT models and a set of human-annotated scores of one of the model’s translation results.

## 2 Evaluation Metrics for Japanese-English Chat Translation

In MT, many traditional baseline metrics remain popular for evaluating MT systems due to their lightweight and fast computation [10]. However, as MT systems improve over time, commonly used metrics struggle to correlate with human judgment at the segment level and fail to evaluate the highest-performing MT systems adequately, thus misleading system development with incorrect conclusions. Among these metrics, BLEU is often considered the de facto standard of MT evaluation metrics. In advanced translation, however, there are many cases in which the description of phrases differs but implies the same meaning. In these cases, BLEU will struggle to perform adequately, even if there is a system that can translate with high quality. It is noted that even though BLEU is fast and can be helpful to assist researchers and developers in quickly performing initial experiments, BLEU should not be the primary evaluation technique in NLP papers [12]. Furthermore, it is argued that, due to its limitations, the common use of BLEU over the past years has negatively affected research decisions in MT [13].

In recent years, many efforts have been conducted to propose better metrics that can measure the quality of MT systems, such as BERTScore [9] and COMET [10], which are proven to correlate better with human judgments. However, only a few provided analysis on their use for evaluating chat translation, which is notably challenging for the Japanese language with its various challenges, such

as the anaphoric pronoun resolution, that occurs more frequently in spoken language than in written language [14]. Furthermore, many works that try to build models for this are evaluated using BLEU, which is often argued as insufficient for advanced models [15]. However, the performance of many recent neural-based metrics has yet to be analyzed. This paper compares and analyses results from different models measured with various metrics and how these metrics correlate with human judgment.

## 3 Experiments

In this section, we explain the data and procedure in the experiment to evaluate the quality of machine-translated sentences using various metrics and provide an analysis of how the result of each metric compares to the human evaluation score and with each other. Firstly, we use three NMT models to translate conversations from Japanese to English and store the translation results. Then, we calculate the score of each model’s translation with various evaluation metrics, from traditional and commonly used metrics like BLEU to more recent neural-based metrics such as BERTScore. Furthermore, we provide an overview of a tool we use for the human evaluation of translated sentence output by the models. Using this evaluation tool, we compile human-annotated scores of machine-translated sentences from one model, measure how each metric correlates with human scores, and share our findings.

### 3.1 Dataset and NMT Model Performance

In our experiment, we use the Business Scene Dialogue (BSD) corpus [16]. We translated the whole development set of the BSD corpus (2,051 chats) using three models, namely our own Transformers-based MT model, M2M100, and MarianMT for Japanese to English language direction. Previously, only BLEU was used to evaluate the results; in this paper, we use six metrics containing both traditional baseline metrics (BLEU, TER, METEOR, and chrF) and recent neural-based metrics (BERTScore and COMET). We can see in Table 1 that even though each metric uses different approaches in evaluating the translated results, all metric seems consistent when deciding which performs the best of the three models. In other words, traditional metrics are still valid if we only want to rank which models

perform better than the others, and the processing time is essential. Table 1 shows that all metrics indicate MarianMT outperforms the other two models.

**Table 1: NMT models performance measured by different metrics**

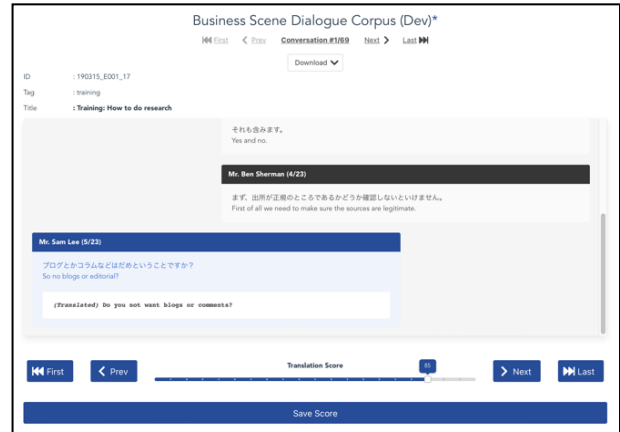
Metric	Ours	M2M100	MarianMT
BLEU	0.14	0.12	<b>0.17</b>
TER	79.70	78.62	<b>75.54</b>
METEOR	0.41	0.39	<b>0.47</b>
chrF	34.31	35.07	<b>39.77</b>
BERTScore	0.92	0.92	<b>0.93</b>
COMET	-0.022	0.026	<b>0.225</b>

### 3.2 Direct Assessment for Human

#### Evaluation and Metrics Performance

Previous works have shown that some metrics are better than others in correlation with human judgment, which is still considered the gold truth of evaluation metrics in machine translation. In this section, we want to explore how these metrics correlate with human judgment when evaluating context-heavy chat sentences from Japanese to English in the BSD corpus. Firstly, to get the human-annotated score, we built a direct assessment tool for the manual evaluation of machine translation. The tool enables the user to view the chat dataset in a familiar chat-like user interface displaying the current segment/chat conversation along with its document-level context (previous chats), then score each translation from 0 (wrong) to 100 (perfect). It is developed following the direct assessment tool used for human evaluation on the WMT 2020 Shared Task on Chat Translation [17]. Figure 1 shows the user interface of this tool.

Using the direct assessment tool for manual evaluation, we collected human-annotated scores of 10 conversations containing 283 chats inside the development set of the BSD corpus, translated by our own MT model. We asked two human annotators to score the translation and averaged the results. Both annotators speak Japanese and English, one is a native English speaker, and the other is a native Japanese speaker. The conversations are picked arbitrarily from the development set.



**Figure 1: User interface of the direct assessment tool**

Table 2 shows how each metric correlates with the human-annotated score for each chat translation. Neural-based metrics pre-trained with language models outperform traditional baseline metrics in correlation with human scores. It is worth noting that neural-based metrics slow down when the computation is done using the CPU only. This can be a limitation when we still experiment with architectures and hyperparameters in the initial phases of training an MT model.

**Table 2: Pearson’s r on each metric**

Metric	Pearson’s r	Elapsed Time
SentenceBLEU	0.2246	<b>1.06 s</b>
TER	-0.1995	1.19 s
METEOR	0.2860	1.23 s
chrF	0.3038	1.07 s
BERTScore (F1)	0.4342	GPU: 6.54 s CPU: 98.73 s
COMET	<b>0.5246</b>	GPU: 16.71 s CPU: 220.71 s

As shown in Table 2, COMET achieves the highest correlation coefficient in our experiment. Based on this, we conduct additional analysis on the results scored by COMET. The aim is to get some clues about how it evaluates translations with a specific language phenomenon, such as the anaphoric zero pronouns, which are prominent, especially in Japanese conversations/chats [18]. Table 3 provides some example sentences where COMET is supposed to give higher scores to the human-translated sentences and lower scores to machine-translated sentences in which the pronouns need to be correctly

translated and could result in a completely different meaning.

**Table 3: Examples where COMET gives lower score to better translations on zero-pronoun sentences**

Example #1		
Source	最後にお話したのはいつでしたっけ？	
Target	When was the last time we talked?	
Machine	When was the last time <b>you</b> talked?	<b>0.931</b>
Human	Do you remember when was the last time we talked?	0.845
Example #2		
Source	できると思います。	
Target	I think so.	
Machine	I think <b>you</b> can.	<b>-0.383</b>
Human	Yes, I think I can.	-0.267

## 4 Conclusion

We analyze various metrics for chat translation from Japanese to English using three models on the BSD developments set. Furthermore, to determine whether the results correlate with human judgment, we compiled a dataset containing ten conversations from the BSD corpus development set (approximately 15% of the total number of exchanges) along with the translation results and their human-annotated score using Direct Assessment. By calculating the Pearson correlation coefficient ( $r$ ), it can be seen that more recent metrics correlate better with human-annotated scores than traditional baseline metrics, with COMET achieving the highest correlation.

There are still, however, limitations on the current best metrics. One major challenge in translating Japanese sentences is when the model needs to decide what pronoun to use when the source sentence does not contain any. Even the best metrics often make simple mistakes in our experiments by giving high scores to translations with wrong pronouns. It is worth further research on improving the metrics to accommodate various language-specific phenomena.

## References

- [1] Läubli, S., Sennrich, R., and Volk, M. 2018. [Has Machine Translation Achieved Human Parity? A Case for Document-level Evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4791–4796, Brussels, Belgium. Association for Computational Linguistics.
- [2] Toral, A., Castilho, S., Hu, K., and Way, A. 2018. [Attaining the Unattainable? Reassessing Claims of Human Parity in Neural Machine Translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.
- [3] Zhang, J., Luan, H., Sun, M., Zhai, F., Xu, J., Zhang, M., and Liu, Y. 2018. [Improving the Transformer Translation Model with Document-Level Context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 533–542, Brussels, Belgium. Association for Computational Linguistics.
- [4] Voita, E., Sennrich, R., and Titov, I. 2019. [When a Good Translation is Wrong in Context: Context-Aware Machine Translation Improves on Deixis, Ellipsis, and Lexical Cohesion](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1198–1212, Florence, Italy. Association for Computational Linguistics.
- [5] Taira, H., Sudoh, K., and Nagata, M. 2012. [Zero Pronoun Resolution can Improve the Quality of J-E Translation](#). In *Proceedings of the Sixth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 111–118, Jeju, Republic of Korea. Association for Computational Linguistics.
- [6] Kudo, T., Ichikawa, H., and Kazawa, H. 2014. [A joint inference of deep case analysis and zero subject generation for Japanese-to-English statistical machine translation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 557–562, Baltimore, Maryland. Association for Computational Linguistics.

- [7] Papineni, K., Roukos, S., Ward, T., and Zhu, W. 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.
- [8] Snover, M., Dorr, B., Schwartz, R., Micciulla, L., and Makhoul, J. 2006. [A Study of Translation Edit Rate with Targeted Human Annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- [9] 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*.
- [10] Rei, R., Stewart, C., Farinha, A.C., and Lavie, A. 2020. [COMET: A Neural Framework for MT Evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- [11] Bexte, M., Horbach, A., and Zesch, T. 2022. [Similarity-Based Content Scoring - How to Make S-BERT Keep Up With BERT](#). In *Proceedings of the 17th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2022)*, pages 118–123, Seattle, Washington. Association for Computational Linguistics.
- [12] Reiter, E. 2018. [A Structured Review of the Validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- [13] Hanna, M. and Bojar, O. 2021. [A Fine-Grained Analysis of BERTScore](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 507–517, Online. Association for Computational Linguistics.
- [14] Riktors, M., Ri, R., Li, T., and Nakazawa, T. 2022. [Japanese–English Conversation Parallel Corpus for Promoting Context-aware Machine Translation Research](#). *Journal of Natural Language Processing*, 28(2): 380-403. <https://doi.org/10.5715/jnlp.28.380>.
- [15] Mathur, N., Baldwin, T., and Cohn, T. 2020. [Tangled up in BLEU: Reevaluating the Evaluation of Automatic Machine Translation Evaluation Metrics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4984–4997, Online. Association for Computational Linguistics.
- [16] Riktors, M., Ri, R., Li, T., and Nakazawa, T. 2019. [Designing the Business Conversation Corpus](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 54–61, Hong Kong, China. Association for Computational Linguistics.
- [17] Farajian, M.A., Lopes, A.V., Martins, A.F.T., Maruf, S., and Haffari, G. 2020. [Findings of the WMT 2020 Shared Task on Chat Translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 65–75, Online. Association for Computational Linguistics.
- [18] Rusli, A. and Shishido, M. 2022. [Zero-Pronoun Annotation Support Tool for the Evaluation of Machine Translation on Conversational Texts](#). *Journal of Natural Language Processing*, 29(2):493-514. <https://doi.org/10.5715/jnlp.29.493>.