

Metric for Automatic Machine Translation Evaluation based on Universal Sentence Representations

Hiroki Shimanaka[†]

Tomoyuki Kajiwara^{†‡}

Mamoru Komachi[†]

[†]Graduate School of Systems Design, Tokyo Metropolitan University, Tokyo, Japan
shimanaka-hiroki@ed.tmu.ac.jp, komachi@tmu.ac.jp

[‡]Institute for Datability Science, Osaka University, Osaka, Japan
kajiwara@ids.osaka-u.ac.jp

Abstract

Sentence representations can capture a wide range of information that cannot be captured by local features based on character or word N-grams. This paper examines the usefulness of universal sentence representations for evaluating the quality of machine translation. Although it is difficult to train sentence representations using small-scale translation datasets with manual evaluation, sentence representations trained from large-scale data in other tasks can improve the automatic evaluation of machine translation. Experimental results of the WMT-2016 dataset show that the proposed method achieves state-of-the-art performance with sentence representation features only.

1 Introduction

This paper describes a segment-level metric for automatic machine translation evaluation (MTE). MTE metrics having a high correlation with human evaluation enable the continuous integration and deployment of a machine translation (MT) system. Various MTE metrics have been proposed in the metrics task of the Workshops on Statistical Machine Translation (WMT) that was started in 2008. However, most MTE metrics are obtained by computing the similarity between an MT hypothesis and a reference translation based on character N-grams or word N-grams, such as SentBLEU (Lin and Och, 2004), which is a smoothed version of BLEU (Papineni et al., 2002), Blend (Ma et al., 2017), MEANT 2.0 (Lo, 2017), and chrF++ (Popović, 2017), which achieved excellent results in the WMT-2017 Metrics task (Bojar et al., 2017). Therefore, they can exploit only limited information for segment-level MTE. In other words, MTE metrics based on character N-grams or word N-grams cannot make full use of sentence representations; they only check for word matches.

We propose a segment-level MTE metric by using universal sentence representations capable of capturing information that cannot be captured by local features based on character or word N-grams. The results of an experiment in segment-level MTE conducted using the datasets for to-English language pairs on WMT-2016 indicated that the proposed regression model using sentence representations achieves the best performance.

The main contributions of the study are summarized below:

- We propose a novel supervised regression model for segment-level MTE based on universal sentence representations.
- We achieved state-of-the-art performance on the WMT-2016 dataset for to-English language pairs without using any complex features and models.

2 Related Work

DPMF_{comb} (Yu et al., 2015a) achieved the best performance in the WMT-2016 Metrics task (Bojar et al., 2016). It incorporates 55 default metrics provided by the Asiya MT evaluation toolkit¹ (Giménez and Màrquez, 2010), as well as three other metrics, namely, DPMF (Yu et al., 2015b), REDp (Yu et al., 2015a), and ENTfP (Yu et al., 2015a), using ranking SVM to train parameters of each metric score. DPMF evaluates the syntactic similarity between an MT hypothesis and a reference translation. REDp evaluates an MT hypothesis based on the dependency tree of the reference translation that comprises both lexical and syntactic information. ENTfP (Yu et al., 2015a) evaluates the fluency of an MT hypothesis.

¹<http://asiya.lsi.upc.edu/>




Figure 1: Outline of Skip-Thought.

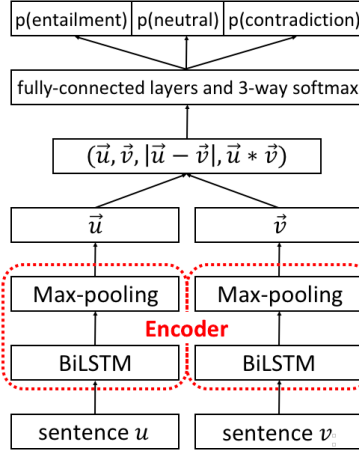


Figure 2: Outline of InferSent.

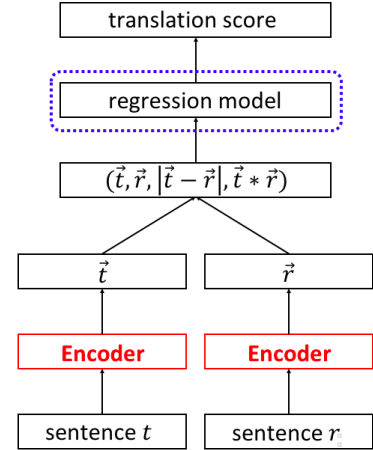


Figure 3: Outline of our metric.

After the success of $\text{DPMF}_{\text{comb}}$, Blend^2 (Ma et al., 2017) achieved the best performance in the WMT-2017 Metrics task (Bojar et al., 2017). Similar to $\text{DPMF}_{\text{comb}}$, Blend is essentially an SVR (RBF kernel) model that uses the scores of various metrics as features. It incorporates 25 lexical metrics provided by the Asiya MT evaluation toolkit, as well as four other metrics, namely, BEER (Stanojević and Sima’an, 2015), CharacTER (Wang et al., 2016), DPMF and ENTFp . BEER (Stanojević and Sima’an, 2015) is a linear model based on character N-grams and replacement trees. CharacTER (Wang et al., 2016) evaluates an MT hypothesis based on character-level edit distance.

$\text{DPMF}_{\text{comb}}$ is trained through relative ranking of human evaluation data in terms of relative ranking (RR). The quality of five MT hypotheses of the same source segment are ranked from 1 to 5 via comparison with the reference translation. In contrast, Blend is trained through direct assessment (DA) of human evaluation data. DA provides the absolute quality scores of hypotheses, by measuring to what extent a hypothesis adequately expresses the meaning of the reference translation. The results of the experiments in segment-level MTE conducted using the datasets for to-English language pairs on WMT-2016 showed that Blend achieved a better performance than $\text{DPMF}_{\text{comb}}$ (Table 2). In this study, we use Blend and propose a supervised regression model trained using DA human evaluation data.

Instead of using local and lexical features,

ReVal^3 (Gupta et al., 2015a,b) proposes using sentence-level features. It is a metric using Tree-LSTM (Tai et al., 2015) for training and capturing the holistic information of sentences. It is trained using datasets of pseudo similarity scores, which is generated by translating RR data, and out-domain datasets of similarity scores of SICK^4 . However, the training dataset used in this metric consists of approximately 21,000 sentences; thus, the learning of Tree-LSTM is unstable and accurate learning is difficult (Table 2). The proposed metric uses sentence representations trained using LSTM as sentence information. Further, we apply universal sentence representations to this task; these representations were trained using large-scale data obtained in other tasks. Therefore, the proposed approach avoids the problem of using a small dataset for training sentence representations.

3 Regression Model for MTE Using Universal Sentence Representations

The proposed metric evaluates MT results with universal sentence representations trained using large-scale data obtained in other tasks. First, we explain two types of sentence representations used in the proposed metric in Section 3.1. Then, we explain the proposed regression model and feature extraction for MTE in Section 3.2.

3.1 Universal Sentence Representations

Several approaches have been proposed to learn sentence representations. These sentence representations are learned through large-scale data so

²<http://github.com/qingsongma/blend>

³<https://github.com/rohitguptacs/ReVal>

⁴<http://clic.cimec.unitn.it/composes/sick.html>

	cs-en	de-en	fi-en	ro-en	ru-en	tr-en
WMT-2015	500	500	500	-	500	-
WMT-2016	560	560	560	560	560	560

Table 1: Number of DA human evaluation datasets for to-English language pairs⁸ in WMT-2015 (Stanojević et al., 2015) and WMT-2016 (Bojar et al., 2016).

that they constitute potentially useful features for MTE. These have been proved effective in various NLP tasks such as document classification and measurement of semantic textual similarity, and we call them universal sentence representations.

First, Skip-Thought⁵ (Kiros et al., 2015) builds an unsupervised model of universal sentence representations trained using three consecutive sentences, such as s_{i-1} , s_i , and s_{i+1} . It is an encoder-decoder model that encodes sentence s_i and predicts previous and next sentences s_{i-1} and s_{i+1} from its sentence representation \vec{s}_i (Figure 1). As a result of training, this encoder can produce sentence representations. Skip-Thought demonstrates high performance, especially when applied to document classification tasks.

Second, InferSent⁶ (Conneau et al., 2017) constructs a supervised model computing universal sentence representations trained using Stanford Natural Language Inference (SNLI) datasets⁷ (Bowman et al., 2015). The Natural Language Inference task is a classification task of sentence pairs with three labels, *entailment*, *contradiction* and *neutral*; thus, InferSent can train sentence representations that are sensitive to differences in meaning. This model encodes sentence pairs u and v and generates features by sentence representations \vec{u} and \vec{v} with a bi-directional LSTM architecture with max pooling (Figure 2). InferSent demonstrates high performance across various document classification and semantic textual similarity tasks.

3.2 Regression Model for MTE

In this paper, we propose a segment-level MTE metric for to-English language pairs. This problem can be treated as a regression problem that estimates translation quality as a real number from an MT hypothesis t and a reference translation r . Once d -dimensional sentence vectors \vec{t} and \vec{r} are generated, the proposed model applies the follow-

ing three matching methods to extract relations between t and r (Figure 3).

- Concatenation: (\vec{t}, \vec{r})
- Element-wise product: $\vec{t} * \vec{r}$
- Absolute element-wise difference: $|\vec{t} - \vec{r}|$

Thus, we perform regression using $4d$ -dimensional features of \vec{t} , \vec{r} , $\vec{t} * \vec{r}$ and $|\vec{t} - \vec{r}|$.

4 Experiments of Segment-Level MTE for To-English Language Pairs

We performed experiments using evaluation datasets of the WMT Metrics task to verify the performance of the proposed metric.

4.1 Setups

Datasets. We used datasets for to-English language pairs from the WMT-2016 Metrics task (Bojar et al., 2016) as summarized in Table 1. Following Ma et al. (2017), we employed all other to-English DA data as training data (4,800 sentences) for testing on each to-English language pair (560 sentences) in WMT-2016.

Features. Publicly available pre-trained sentence representations such as Skip-Thought⁵ and InferSent⁶ were used as the features mentioned in Section 3. Skip-Thought is a collection of 4,800-dimensional sentence representations trained on 74 million sentences of the BookCorpus dataset (Zhu et al., 2015). InferSent is a collection of 4,096-dimensional sentence representations trained on both 560,000 sentences of the SNLI dataset (Bowman et al., 2015) and 433,000 sentences of the MultiNLI dataset (Williams et al., 2017).

Model. Our regression model used SVR with the RBF kernel from scikit-learn⁹. Hyperparameters were determined through 10-fold cross

⁵<https://github.com/ryankiros/skip-thoughts>

⁶<https://github.com/facebookresearch/InferSent>

⁷<https://nlp.stanford.edu/projects/snli/>

⁸en: English, cs: Czech, de: German, fi: Finnish, ro: Romanian, ru: Russian, tr: Turkish

⁹<http://scikit-learn.org/stable/>

	cs-en	de-en	fi-en	ro-en	ru-en	tr-en	Avg.
SentBLEU (Bojar et al., 2016)	0.557	0.448	0.484	0.499	0.502	0.532	0.504
Blend (Ma et al., 2017)	0.709	0.601	0.584	0.636	0.633	0.675	0.640
DPMF _{comb} (Bojar et al., 2016)	0.713	0.584	0.598	0.627	0.615	0.663	0.633
ReVal (Bojar et al., 2016)	0.577	0.528	0.471	0.547	0.528	0.531	0.530
SVR with Skip-Thought	0.665	0.571	0.609	0.677	0.608	0.599	0.622
SVR with InferSent	0.679	0.604	0.617	0.640	0.644	0.630	0.636
SVR with InferSent + Skip-Thought	0.686	0.611	0.633	0.660	0.649	0.646	0.648

Table 2: Segment-level Pearson correlation of metric scores and DA human evaluations scores for to-English language pairs in WMT-2016 (newstest2016).

validation using the training data. We examined all combinations of hyper-parameters among $C \in \{0.01, 0.1, 1.0, 10\}$, $\epsilon \in \{0.01, 0.1, 1.0, 10\}$, and $\gamma \in \{0.01, 0.1, 1.0, 10\}$.

There are three comparison methods: Blend (Ma et al., 2017), DPMF_{comb} (Yu et al., 2015a), and ReVal (Gupta et al., 2015a,b), as described in Section 2. Blend and DPMF_{comb} are MTE metrics that exhibited the best performance in the WMT-2017 Metrics task (Bojar et al., 2017) and WMT-2016 Metrics task, respectively. We compared the Pearson correlation of each metric score and DA human evaluation scores.

4.2 Result

As can be seen in Table 2, the proposed metric, which combines InferSent and Skip-Thought representations, surpasses the best performance in three out of six to-English languages pairs and achieves state-of-the-art performance on average.

4.3 Discussion

These results indicate that it is possible to adopt universal sentence representations in MTE by training a regression model using DA human evaluation data. Since Blend is an ensemble method using combinations of various MTE metrics as features, our results show that universal sentence representations can consider information more abundantly than a complex model. Since ReVal is also based on sentence representations, we conclude that universal sentence representations trained on a large-scale dataset are more effective for MTE tasks than sentence representations trained on a small or limited in-domain dataset.

4.4 Error Analysis

We re-implemented Blend¹⁰ (Ma et al., 2017) and compared the evaluation results with the proposed metric.¹¹

We analyzed 20% of the pairs of MT hypotheses and reference translations (112 sentence pairs \times 6 languages = 672 sentence pairs) in descending order of DA human score in each language pair. In other words, the top 20% of MT hypotheses that were close to the meaning of the reference translations for each language pair were analyzed. Among these, only Blend estimates the translation quality as high for 70 sentence pairs, and only our metric estimates the translation quality as high for 88 sentence pairs.

Surface. Among pairs estimated to have high translation quality by each method, there were 26 pairs in Blend and 42 pairs in the proposed method with a low word surface matching rate between MT hypotheses and reference translations. This result shows that the proposed metric can evaluate a wide range of sentence information that cannot be captured by Blend.

Unknown words. There were 26 MT hypotheses consisting of words that were treated as unknown words in Skip-Thought or InferSent that were correctly evaluated in Blend. On the other hand, there were 26 MT hypotheses that were correctly evaluated in the proposed metric. This result shows that the proposed metric is affected by unknown words. However, it is also true that there are some MT hypotheses containing unknown words that can be correctly evaluated.

¹⁰<http://github.com/qingsongma/blend>

¹¹The average Pearson correlation of all language pairs after re-implementing Blend was 0.636, which is a little lower than the value reported in their paper. However, we judged that the following discussion will not be affected by this difference.

Therefore, we analyzed further by focusing on sentence length. There were 17 MT hypotheses consisting of words that were treated as unknown words by either Skip-Thought or InferSent with a short length (15 words or less) that were correctly evaluated in Blend. However, in the proposed metric, there were only two MT hypotheses that were correctly evaluated. This result indicates that the shorter the sentence, the more likely is the proposed metric to be affected by unknown words.

5 Conclusions

In this study, we tried to apply universal sentence representation to MTE based on the DA of human evaluation data. Our segment-level MTE metric achieved the best performance on the WMT-2016 dataset. We conclude that:

- Universal sentence representations can consider information more comprehensively than an ensemble metric using combinations of various MTE metrics based on features of character or word N-grams.
- Universal sentence representations trained on a large-scale dataset are more effective than sentence representations trained on a small or limited in-domain dataset.
- Although a metric based on SVR with universal sentence representations is not good at handling unknown words, it correctly estimates the translation quality of MT hypotheses with a low word matching rate with reference translations.

Following the success of InferSent (Conneau et al., 2017), many works (Wieting and Gimpel, 2017; Cer et al., 2018; Subramanian et al., 2018) on universal sentence representations have been published. Based on the results of our work, we expect that the MTE metric will be further improved using these better universal sentence representations.

References

- Ondřej Bojar, Yvette Graham, and Amir Kamran. 2017. Results of the WMT17 Metrics Shared Task. In *Proceedings of the Second Conference on Machine Translation*, pages 489–513.
- Ondřej Bojar, Yvette Graham, Amir Kamran, and Miloš Stanojević. 2016. Results of the WMT16 Metrics Shared Task. In *Proceedings of the First Conference on Machine Translation*, pages 199–231.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A Large Annotated Corpus for Learning Natural Language Inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642.
- Daniel Cer, Yinfei Yang, Sheng yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Universal Sentence Encoder. *arXiv preprint arXiv:1803.11175*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680.
- Jesús Giménez and Lluís Màrquez. 2010. Asiya: An Open Toolkit for Automatic Machine Translation (Meta-) Evaluation. *The Prague Bulletin of Mathematical Linguistics*, (94):77–86.
- Rohit Gupta, Constantin Orasan, and Josef van Genabith. 2015a. ReVal: A Simple and Effective Machine Translation Evaluation Metric Based on Recurrent Neural Networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1066–1072.
- Rohit Gupta, Constantin Orasan, and Josef van Genabith. 2015b. Machine Translation Evaluation using Recurrent Neural Networks. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 380–384.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-Thought Vectors. In *Advances in Neural Information Processing Systems* 28, pages 3294–3302.
- Chin-Yew Lin and Franz Josef Och. 2004. ORANGE: a Method for Evaluating Automatic Evaluation Metrics for Machine Translation. In *Proceedings of the 20th International Conference on Computational Linguistics*, pages 501–507.
- Chi-Kiu Lo. 2017. MEANT 2.0: Accurate semantic MT evaluation for any output language. In *Proceedings of the Second Conference on Machine Translation*, pages 589–597.
- Qingsong Ma, Yvette Graham, Shugen Wang, and Qun Liu. 2017. Blend: a Novel Combined MT Metric Based on Direct Assessment —CASICT-DCU submission to WMT17 Metrics Task. In *Proceedings of the Second Conference on Machine Translation*, pages 598–603.

- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. BLEU: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the Second Conference on Machine Translation*, pages 612–618.
- Miloš Stanojević, Philipp Koehn, and Ondřej Bojar. 2015. Results of the WMT15 Metrics Shared Task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 256–273.
- Miloš Stanojević and Khalil Sima'an. 2015. BEER 1.1: ILLC UvA submission to metrics and tuning task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 396–401.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. Learning General Purpose Distributed Sentence Representations via Large Scale Multi-task Learning. In *Proceedings of the 6th International Conference on Learning Representations*, pages 1–16.
- Weiyue Wang, Jan-Thorsten Peter, Hendrik Rosendahl, and Hermann Ney. 2016. Character: Translation Edit Rate on Character Level. In *Proceedings of the First Conference on Machine Translation*, pages 505–510.
- John Wieting and Kevin Gimpel. 2017. Pushing the Limits of Paraphrastic Sentence Embeddings with Millions of Machine Translations. *arXiv preprint arXiv:1711.05732*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. *arXiv preprint arXiv:1704.05426*.
- Hui Yu, Qingsong Ma, Xiaofeng Wu, and Qun Liu. 2015a. CASICT-DCU Participation in WMT2015 Metrics Task. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 417–421.
- Hui Yu, Xiaofeng Wu, Wenbin Jiang, Qun Liu, and Shouxun Lin. 2015b. An Automatic Machine Translation Evaluation Metric Based on Dependency Parsing Model. *arXiv preprint arXiv:1508.01996*.
- Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books. *arXiv preprint arXiv:1506.06724*.