Estimating the Quality of Translated Medical Texts using Back Translation & Resource Description Framework

Vinay Neekhra*, Dipti Misra Sharma

Language Technology Research Center, Kohli Center on Intelligent Systems,
International Institute of Information Technology, Hyderabad (IIIT-Hyderabad)

Abstract

How can we effectively estimate the quality of translated texts in the medical field, where back-translation is usually available and/or recommended for sensitive documents. This paper proposes a novel metric, GATE[1], for translation quality estimation task, leveraging the Resource Description Framework (RDF) to encode both semantic and syntactical information of the original and back-translated sentences into RDF graphs. The distance between these graphs is measured to get the semantic similarity score to assess the quality of the translation. Unlike traditional metrics like BLEU and METEOR, our approach is reference-less, capturing both semantic and syntactical information for a comprehensive assessment of translation quality. Our results correlate better with human judgment, giving a better Pearson correlation (0.357) as compared to BLEU (0.200), thereby showing ~70% improvement over BLEU. Our research shows that, in the field of translation evaluation, existing resources like back-translation and RDF could be useful.

Keywords

Translation Quality Estimation, Resource Description Framework (RDF), Back Translation, GATE

1. Introduction

A drug trial in the medical domain incorporates a mandatory consent form called a Medical Consent Form (MCF), which informs the patient about the experiment and its potential side effects. There is a legal requirement for the MCF to be in the patient's mother tongue and for it to be easy to understand. A human translator translates the original MCF into the patient's mother tongue. As MCFs are sensitive documents, evaluating the quality of translated texts is crucial to ensure faithfulness to the original texts (see Section 1.1 for an example).

One way to evaluate the quality of the translated texts is using back-translation (see Section 3.1), wherein the translated text is translated back into the original language. The original and back-translated texts are then compared to estimate the quality of the translation. Back-translation is a prominent way to assess the quality of translated texts in domains, such as medical documents, where accuracy and precision are paramount [1][2].

Experienced professionals are responsible for carrying out all three procedures (see Figure 1), namely: initial translation from the source language to the target language, followed by translation from the target language back to the source language, and ultimately, comparison between the original text and the back-translated texts. Our efforts are focused on reducing the efforts of human evaluators comparing the original and back-translated texts by automating the task of evaluating the quality of translated texts.

While human evaluation has traditionally served as a benchmark for assessing translation quality, it is often expensive, time-consuming, and subjective. As an alternative, automatic evaluation metrics such as BLEU[3], METEOR[4], etc., have been developed to provide a more efficient and objective means of evaluation, with BLEU being the most commonly used metric (see Section 2 for related work). This field of research, called translation quality estimation (QE), is an area of research concerned with

---

evaluating the quality of translated texts when gold standard translations (called reference texts) are unavailable.

In this paper, we propose a novel translation evaluation metric, GATE (Graphical Assessment for Translation quality Estimation), which leverages back-translation (see Section 3.1) and the Resource Description Framework (RDF) (see Section 3.2). GATE encodes both semantic and syntactical information of the original and back-translated sentences into RDF graphs, allowing for a reference-less, semantically-aware assessment of translation quality.

For sensitive documents in the medical field, such as medical consent forms and qualitative research, back-translation is a common practice to ensure the faithfulness of translations [1][2]. GATE capitalizes on this by integrating back-translation into its evaluation framework, providing a comprehensive and reliable assessment of translation quality. To estimate the quality of translated texts, we encode the meaning of these sentences into graphs using the Resource Description Framework (RDF) and then compare these graphs to come up with a similarity score (See Figure 4). GATE shows a higher correlation (0.357) with human judgment than BLEU (0.200). (see Section 4 for the experiment details). In the next Section 1.1, we discuss the significance of translation evaluation, highlighting the context and motivation behind our research efforts.

1.1. Significance of Translation Evaluation

Consider the following sentence from a medical consent form for a vaccine trial, translated to the patient's mother tongue *(Tamil language)* where the original consent form is in English.

- *Source text:*   There are no side effects mentioned previously.

To comply with legal requirements, the consent form was translated into Tamil by hospital authorities, resulting in two translated versions. For evaluating the translation quality, the translated MCF was back-translated to English, yielding the following results:

- *Back Translation 1:*   No side effects which were mentioned previously
- *Back Translation 2:*   It has already been mentioned that it does not have any side-effects

As seen above, the first back-translated sentence is semantically similar to the source text and preserves the original intent. The second back translated text, on the other hand, conveys that —*as previously mentioned, there are no side-effects*—, whereas the original intent was that no side-effects have been observed yet, thus raising ethical and legal concerns.

Thus, it is crucial, that translated texts are evaluated for their faithfulness to the original text, especially in the medical domain. In the next subsection, we highlight the contributions of our work.

1.2. Contributions

1. This paper presents a novel approach, GATE, for translation quality estimation task by utilizing back-translation and leveraging knowledge graphs (namely, Resource Description Framework) for encoding the meaning of original and back-translated texts to come up with a translation quality estimation score.
2. GATE incorporates both syntactic and semantic information, leading to improved evaluation scores. Our approach is applicable to both machine-translated and human-translated texts. Our experiments demonstrate a better correlation with human judgment compared to BLEU, with a Pearson correlation of 0.357 compared to the most commonly used metric, BLEU's 0.200.
3. Our approach eliminates the need for reference texts by comparing the source text directly with its back-translated counterpart. This makes our approach reference-less and thus valuable for

scenarios where reference texts are not available for translation evaluation (such as medical consent forms).

4. While our results do not surpass the current state-of-the-art, our metric, GATE, offers distinct advantages such as requiring no training, being computationally lightweight, being available for low-resource languages, and operating without the need for extensive training data, unlike neural network-based methods like COMET [5].

The paper is structured as follows: Section 2 reviews related work in the area of translation evaluation, discussing the limitations of existing metrics. Section 3 builds the foundation of our work, providing an overview of back-translation along with its significance, introduces Knowledge Graphs in general, and describes Resource Description Framework (RDF) and FRED RDF graphs. Section 4 details the experiment design and methodology leading to the creation of GATE. The results of our experiments are presented in Section 5, along with a discussion of the insights gained from our research efforts while also addressing the current limitations of our metric. Finally, Section 6 and Section 7 conclude the paper along with outlining the directions for future research.

2. Related work

Existing metrics for translation evaluation, such as BLEU[3], METEOR[4], NIST[6], and TER[7], have been widely utilized in the field, with BLEU being the most commonly used among them. BLEU compares the translated sentence with a reference sentence. It operates on word group matching using an n-gram model and remains popular due to its simplicity. In contrast, METEOR was developed as a successor to BLEU to account for synonyms and other variations in language. Usually, the quality of translation is evaluated at the sentence level, but word and document level QE are also possible [8].

However, these metrics have inherent limitations. Many traditional metrics are categorized as n-gram matching metrics, relying on handcrafted features to estimate translation quality by counting the number and fraction of n-grams shared between a candidate translation hypothesis and one or more human references. This restricts their ability to capture nuanced meaning, particularly in complex and domain-specific texts. They often rely on surface-level similarity measures and may necessitate reference translations, typically provided by humans as a standard of perfection.

More recent approaches have explored the use of word embeddings as an alternative to n-gram matching for capturing word semantic similarity. Metrics like BLEU2VEC[9], BERT SCORE[10], and COMET[5] create alignments between reference and hypothesis segments in an embedding space to compute a score reflecting semantic similarity. COMET, a notable metric in this domain, has demonstrated remarkable results for translation evaluation. However, to train these models, the availability of word embeddings for low-resource languages remains a significant challenge.

However, these metrics may still need to catch up in capturing the full range of nuances captured by human judgments. Challenges with existing metrics include their reliance on reference texts for comparison, requiring semantic exactness at the word level, susceptibility to differences in lexical structure (such as word order), and the tendency to measure semantic relatedness rather than semantic similarity, huge data requirement for training models thus not well-suited for low-resource languages.

3. Preliminaries

This section lays out the foundation required for our experiment design.

3.1. Back Translation:

Back translation is a process where a translated text is translated back into the original language (source language) by a different translator [11]. In Figure 1, translation and back-translation processes between English and French are illustrated, as depicted by [12].
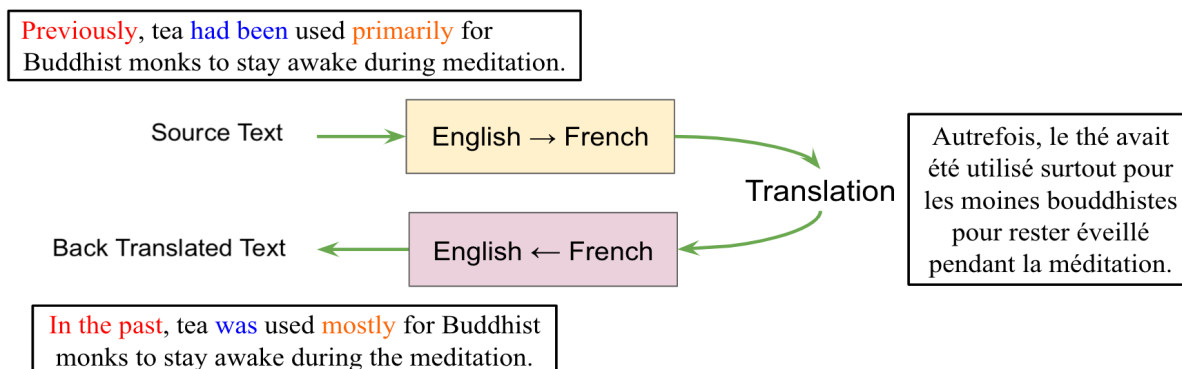
Figure 1: Example of Back Translation (best viewed in color)

Back translation is recommended in the domains where the content subjected to translation is too sensitive and needs to be double-checked. The back-translation method is widely used in medical research and clinical trials, as it is required by Ethics Committees and regulatory authorities in several countries [1]. This allows us to compare the back-translated text with the original text to evaluate the quality of the translation.

The rationale behind using back-translation is that for sensitive documents in the medical domain, back-translation is a recommended practice to cross-verify that the translation adheres to the intended meaning. Usually, back-translation is mandatory in case of quality assessment of medical consent forms, so this is not an overhead in this particular scenario and is generally recommended for medical, legal, market research, and government agencies working in public health, safety, and legal matters. We are utilizing this for translation evaluation. We aim to address the specific needs of these domains to ensure the faithfulness of the translated texts. Our efforts are to use already available back-translation texts for the translation evaluation tasks.

3.2. Resource Description Framework

The Resource Description Framework (RDF) is a W3C standard for data representation on the Web. RDF provides a foundation for encoding information in a structured way for the Semantic Web [13]. It is particularly useful for representing knowledge about entities and the relationships between them.

3.2.1. Components of RDF

RDF consists of triplets, which are fundamental units of information. These triplets, also known as RDF triples, form the building blocks for representing knowledge within an RDF graph. Each RDF triple is composed of three elements:

1. **Subject:** The resource (entity) being described. (e.g., "The patient")
2. **Predicate:** The property or characteristic of the subject, denoted by directed arrows. (e.g., "has diagnosisof")
3. **Object:** The value associated with the predicate for the subject. (e.g., "pneumonia")


Figure 2: RDF Triple for the sentence "The patient has diagnosis of pneumonia"

In Figure 2, the RDF triple depicts a statement about a patient having a diagnosis of pneumonia. In the context of our research, we leverage RDF to capture the semantics of the sentences, enabling a more nuanced evaluation of translation quality compared to traditional metrics.

3.2.2. FRED RDF Graphs

Our research is based on RDF graphs provided by FRED (Framework for RDF-based Extraction and Disambiguation) [14] to capture semantic nuances in translated texts. At its core, FRED leverages the Resource Description Framework (RDF) to construct semantic graphs that capture the relationships and entities present in the text. FRED bridges the gap between unstructured text and structured knowledge representation, employing Semantic Web technologies to extract and disambiguate information from textual data. Figure 3 shows the RDF graph for the sentence "An experimental drug is one which has not been approved by FDA.".
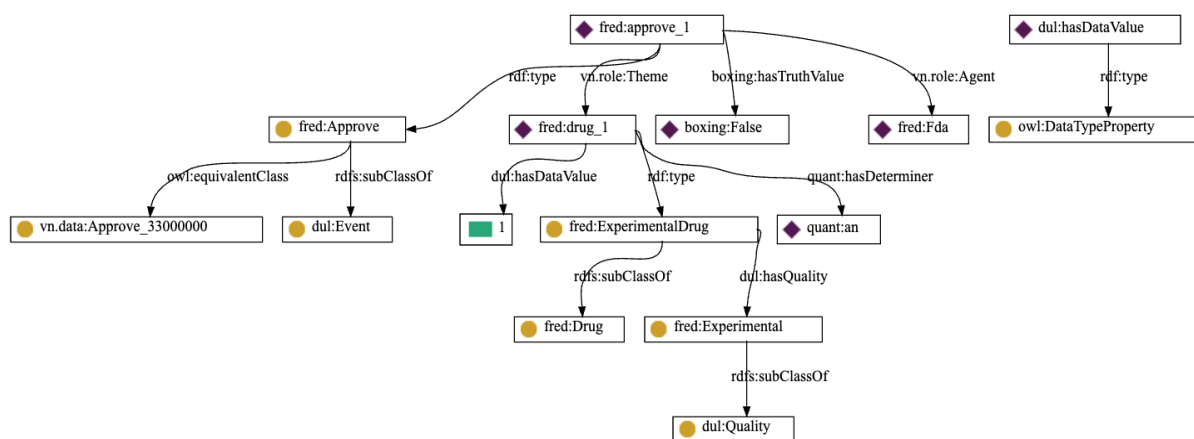
Figure 3: FRED RDF graph for "An experimental drug is one which has not been approved by FDA." taken from a medical consent form.

4. Experiment Design

We conduct a comparative experiment to evaluate the efficacy of our proposed RDF-based evaluation metric, GATE, in comparison to the baseline metric BLEU and its correlation with human judgment. To obtain baseline BLEU scores, we are using iBLEU [15]. The evaluation procedure, outlined in Algorithm 1, explains the comparison of RDF graphs generated through the FRED API, which can be accessed at http://wit.istc.cnr.it/stlab-tools/fred/demo/.
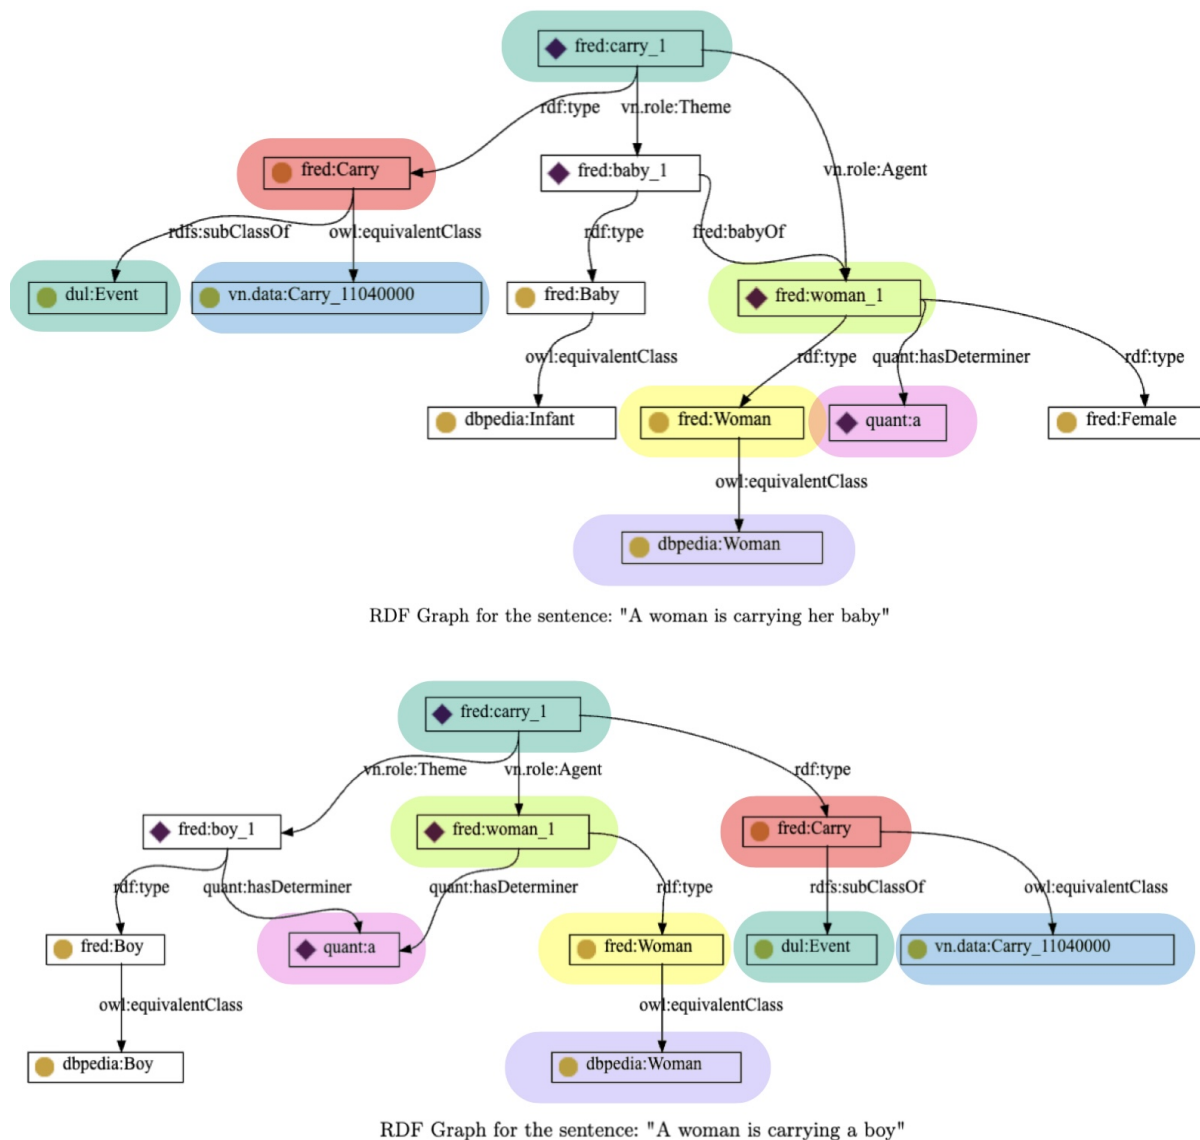
4.1. Dataset

Our experiments were done on the selected medical consent forms and the sentences from Semantic Textual Similarity (STS) Benchmark Dataset [16] to evaluate the effectiveness of GATE in capturing semantic similarity compared to BLEU. The medical consent forms dataset has around 250 original sentence, their corresponding translations, and the back-translated texts, all provided by human translators. Due to the selected availability of medical data, we augmented our analysis with the STS benchmark dataset. In total, our experiments were conducted on 500 sentence pairs, with 250 pairs sourced from medical consent forms provided by a medical institute.

4.2. Graph comparison and GATE Score

We are comparing the source sentence with the back-translated text by constructing RDF graphs for both. The distance between graphs is measured as the Jaccard similarity coefficient [17] between the entities in the graphs. This way, the distance between the source and the back-translated sentence graph is normalized between 0 and 1, where 1 denotes an exact match, and 0 denotes no similarity. Algorithm 1 outlines the steps in the evaluation process. Specifically, for source sentence $\mathbf{s}_k$, and the back-translated text $\mathbf{b}_k$, the GATE Score is calculated as follows:

$$G_k = \frac{entities(\mathbf{s}_k) \cap entities(\mathbf{b}_k)}{entities(\mathbf{s}_k) \cup entities(\mathbf{b}_k)}$$

RDF Graph for the sentence: "A woman is carrying her baby"

RDF Graph for the sentence: "A woman is carrying a boy"

Figure 4: Graph Comparison for measuring semantic similarity. Common nodes are highlighted in multiple colors. In these two graphs there are 8 common nodes, and total unique nodes are 15. (best viewed in color)

For the Figure 4, the GATE Score is calculated as:

$$G = \frac{8 \text{ (number of common entities)}}{15 \text{ (total unique nodes in both the graphs)}} = 0.53$$

Algorithm 1 : GATE Score evaluation process	

Require: All source sentences $\mathbf{s}_k \in \mathbf{S}$ and target sentences $\mathbf{t}_k \in \mathbf{T}$ of n sentence pairs
Ensure: sentence-level scores $\mathbf{G}_k$

1: **for** each sentence pair $\{\mathbf{s}_k, \mathbf{t}_k\} \in \{\mathbf{S}, \mathbf{T}\}$ **do**
2:      $\mathbf{b}_k \leftarrow$ back-translation of $\mathbf{t}_k$ (either already available or obtained using Google Translate)
3:
4:      $entities(\mathbf{s}_k) \leftarrow$ RDF graph nodes of $\mathbf{s}_k$ using FRED
5:      $entities(\mathbf{b}_k) \leftarrow$ RDF graph nodes of $\mathbf{b}_k$ using FRED
6:
7:      **common** $\leftarrow \{x \mid x \in entities(\mathbf{s}_k)$ **and** $x \in entities(\mathbf{b}_k)\}$
8:      **unison** $\leftarrow \{x \mid x \in entities(\mathbf{s}_k)$ **or** $x \in entities(\mathbf{b}_k)\}$
9:
10:      $\mathbf{G}_k \leftarrow \dfrac{common}{unison}$
11:
12: **end for**

In the next section, we present the findings of our experiments along with a discussion of the insights gained from our research efforts while also addressing the current limitations of our metric.

5. Results & Discussion

Our experiment implemented the proposed GATE metric alongside the baseline metric, BLEU. We calculated the Pearson correlation between the BLEU score and GATE score against human judgment on the experiment dataset. Our results in Table 1, show that GATE achieves a significantly higher correlation with human judgment in translation evaluation tasks compared to the widely used metric, BLEU. Specifically, GATE exhibits a ~70% improvement in correlation on the experiment data, with a Pearson correlation coefficient of 0.357 compared to BLEU's 0.200. The higher correlation underscores the effectiveness of leveraging RDF graphs in capturing semantic information, thereby improvement in correlation with human judgments.

Table 1
System-wide Pearson correlation of BLEU and GATE with human judgments on MCFs Data and STS Benchmark Dataset

Metric	Pearson Correlation
BLEU	0.200
GATE	**0.357**

Table 2 shows examples with corresponding human evaluation scores, GATE scores, and BLEU scores. These examples serve to highlight GATE's capability to better reflect human perception of semantic similarity, as evidenced by its closer alignment with human judgments compared to BLEU scores. In summary, our findings indicate that integrating RDF graphs with already existing back-translated texts holds promise for reference-free translation evaluation. This metric can potentially assist human evaluators who evaluate the translation of sensitive documents using back-translated texts.

Using RDF for translation evaluation could be helpful as they 'encode' real-world semantics akin to how embeddings work in neural network frameworks (such as COMET), contrasting with metrics that are based on lexical level information for translation evaluation (such as BLEU). This work has the potential to pave the way for utilizing knowledge graphs in the field of translation evaluation alongside existing resources, such as word embeddings and LLM-based frameworks. Our experiments reinforce our belief, demonstrating that using knowledge graphs to encode meaning is helpful and gets better results than the baseline metrics.

Table 2

GATE vs. BLEU score against human evaluation. Selected examples from the experiment run on STS dataset. Higher correlation with human judgment are marked in bold.

Serial	Hypothesis	Reference	Human	GATE	BLEU
1.	A man is erasing a chalk board	The man is erasing the chalk board	*1.00*	**0.65**	0.60
2.	Three men are playing guitars	Three men on stage are playing guitars	*0.75*	0.45	**0.60**
3.	A woman is carrying a boy	A woman is carrying her baby	*0.47*	**0.53**	0.63
4.	A woman peels a potato	A woman is peeling a potato.	*1.00*	**1.00**	0.52

Given that RDF is currently available only in English and our metric compares graphs of original and back-translated texts for translation evaluation, our metric is presently only applicable where English is the source language. However, the target language can be any other language as long as back-translation is available.

While our results do not surpass state-of-the-art performance, they serve as a proof-of-concept, showcasing the effectiveness of leveraging RDF graphs for translation evaluation tasks. As FRED accommodates large sentences as well, our future work will involve working with more extensive real-world translated medical data and testing our methodology on larger sentences to demonstrate its effectiveness comprehensively. These results underscore the advantages of GATE over traditional metrics like BLEU and motivate further validation of GATE's applicability on real-world data particularly in domains like medicine, along with continuing our exploration for further improvement of the metric.

6. Conclusion

In this paper, we introduce GATE, a novel metric based on the Resource Description Framework (RDF) designed for assessing the quality of translated medical texts for which back-translation is available. To showcase the effectiveness of our metric, we conducted experiments using selected medical data and the STS benchmark dataset, comparing the results against the baseline metric, BLEU, and human judgment scores. Notably, GATE exhibits a stronger correlation with human judgment than BLEU, achieving a higher Pearson correlation coefficient (0.357 compared to BLEU's 0.200), representing approximately a ~70% improvement over BLEU, the most commonly used metric.

By leveraging back-translation and using RDF graphs to encode both semantic and syntactical information, GATE provides a reference-less and semantically aware assessment of translation quality. In comparison with the more advanced Large Language Model (LLM)-based metrics such as COMET, our metric is computationally much lighter. It works for any target language, including low-resource languages, and does not require any data training. Our research shows that, in the field of translation evaluation, existing resources like back-translation and Resource Description Framework could be helpful in real-world scenarios such as the medical domain.

7. Future Directions

As part of future work, we would like to explore:

1. Conducting further experiments to validate the efficacy of GATE on real-world translated medical data.

2. Since Translation and Summarization can both be viewed as natural language generation from a textual context, we aim to explore knowledge graphs such as RDF in the area of evaluating summarization or similar natural language generation tasks. Investigate the utilization of knowledge graphs for tasks beyond translation evaluation, such as summarization.

3. For calculating GATE score, experimenting with different formulas incorporating variations in weights of entities, incoming edges, and outgoing edges.

4. Addressing the challenge of language dependency in GATE by incorporating multilingual knowledge graphs since FRED works only with English texts. A primary avenue for future work, will be looking into the inclusion of other knowledge graphs available in other languages, making GATE language independent.

5. Development of a software similar to iBLEU for integrating FRED API to facilitate automatic scoring of source and back-translated texts, enhanced visualization, and accessibility of the RDF metric.

6. [18] shows that back-translation could be useful for improving the translation quality for low-resource languages. Our future work is to combine neural networks with back-translation and knowledge graphs in the area of translation evaluation for low-resource languages. Our future work aims to combine these technologies along with knowledge graphs (such as Knowledge Graph Embeddings) to improve our metric, making it suitable for evaluating translated sensitive texts and investigating the potential of combining neural networks with back-translation and knowledge graphs to improve translation quality, particularly for low-resource languages.

Acknowledgments

References

[1] D. Grunwald, N. Goldfarb, Back translation for quality control of informed consent forms 2 (2006).

[2] H.-Y. Chen, J. R. Boore, Translation and back-translation in qualitative nursing research: methodological review, Journal of Clinical Nursing 19 (2010) 234–239.

[3] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, BLEU: a method for automatic evaluation of machine translation, Association for Computational Linguistics, 2002, pp. 311–318. doi:10.3115/1073083.1073135.

[4] S. Banerjee, A. Lavie, METEOR: An automatic metric for MT evaluation with improved correlation with human judgments, in: Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization, Association for Computational Linguistics, 2005, pp. 65–72. URL: https://aclanthology.org/W05-0909.

[5] R. Rei, J. G. C. de Souza, D. Alves, C. Zerva, A. C. Farinha, T. Glushkova, A. Lavie, L. Coheur, A. F. T. Martins, COMET-22: Unbabel-IST 2022 submission for the metrics shared task, in: Proceedings of the Seventh Conference on Machine Translation (WMT), Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Hybrid), 2022, pp. 578–585. URL: https://aclanthology.org/2022.wmt-1.52.

[6] G. Doddington, Automatic evaluation of machine translation quality using n-gram co-occurrence statistics, in: Proceedings of the second international conference on Human Language Technology Research, 2002, pp. 138–145.

[7] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, J. Makhoul, 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, 2006, pp. 223–231.

[8] L. Specia, C. Scarton, G. H. Paetzold, G. Hirst, Quality estimation for machine translation, volume 11, Springer, 2018.

[9] A. Tättar, M. Fishel, bleu2vec: the painfully familiar metric on continuous vector space steroids, in: O. Bojar, C. Buck, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, M. Huck, A. J. Yepes, P. Koehn, J. Kreutzer (Eds.), Proceedings of the Second Conference on Machine Translation,

Association for Computational Linguistics, Copenhagen, Denmark, 2017, pp. 619–622. URL: https://aclanthology.org/W17-4771. doi:10.18653/v1/W17-4771.

[10] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, Y. Artzi, Bertscore: Evaluating text generation with bert, arXiv preprint arXiv:1904.09675 (2019).

[11] A. Q, What is back translation?, 2021. URL: https://gtelocalize.com/what-is-back-translation/.

[12] T. H. Trinh, T. Le, P. Hoang, M. Luong, A tutorial on data augmentation by backtranslation, https://github.com/vietai/dab (2019).

[13] World Wide Web Consortium, Resource description framework (rdf) syntax specification (revised), 1998. URL: https://www.w3.org/TR/PR-rdf-syntax/.

[14] A. Gangemi, V. Presutti, D. R. Recupero, A. G. Nuzzolese, F. Draicchio, M. Mongiovà, Semantic Web Machine Reading with FRED, Semantic Web 8 (2017) 873–893.

[15] N. Madnani, ibleu: Interactively debugging and scoring statistical machine translation systems, in: 2011 IEEE Fifth International Conference on Semantic Computing, 2011, pp. 213–214. doi:10.1109/ICSC.2011.36.

[16] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, L. Specia, Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation, in: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), Association for Computational Linguistics, 2017. URL: http://dx.doi.org/10.18653/v1/S17-2001. doi:10.18653/v1/s17-2001.

[17] Wikipedia contributors, Jaccard similarity, https://en.wikipedia.org/wiki/Jaccard_index, 2023.

[18] I. Abdulmumin, B. S. Galadanci, A. Isa, Enhanced back-translation for low resource neural machine translation using self-training, in: S. Misra, B. Muhammad-Bello (Eds.), Information and Communication Technology and Applications, Springer International Publishing, Cham, 2021, pp. 355–371.