

Measuring Uncertainty in Translation Quality Evaluation (TQE)

Serge Gladkoff¹ Irina Sorokina^{1,2} Lifeng Han^{3,5*} Alexandra Alekseeva^{4,‡}

¹ Logrus Global LLC, ² Tver State University, ³ The University of Manchester ⁴ ROKO Labs (‡ deceased)

⁵ ADAPT Research Centre, Dublin City University

{serge.gladkoff, irina.sorokina}@logrusglobal.com

* *corresponding*: lifeng.han@{manchester.ac.uk, adaptcentre.ie}

Abstract

From both human translators (HT) and machine translation (MT) researchers' point of view, translation quality evaluation (TQE) is an essential task. Translation service providers (TSPs) have to deliver large volumes of translations which meet customer specifications with harsh constraints of required quality level in tight time-frames and costs. MT researchers strive to make their models better, which also requires reliable quality evaluation. While automatic machine translation evaluation (MTE) metrics and quality estimation (QE) tools are widely available and easy to access, existing automated tools are not good enough, and human assessment from professional translators (HAPs) are often chosen as the golden standard (Han et al., 2021b). Human evaluations, however, are often accused of having low reliability and agreement. Is this caused by subjectivity or statistics is at play? How to avoid the entire text to be checked and be more efficient with TQE from cost and efficiency perspectives, and what is the optimal sample size of the translated text, so as to reliably estimate the translation quality of the entire material? This work carries out such a motivated research to correctly estimate the confidence intervals (Brown et al., 2001) depending on the sample size of translated text, e.g. the amount of words or sentences, that needs to be processed on TQE workflow step for confident and reliable evaluation of overall translation quality. The methodology we applied for this work is from Bernoulli Statistical Distribution Modeling (BSDM) and Monte Carlo Sampling Analysis (MCSA).

Keywords: Translation Quality Evaluation, Quality Estimation, Post-editing Distance, Confidence Intervals, Monte Carlo Modeling, Bernoulli Statistics

1. Introduction

Machine Translation (MT) is one of the pioneering artificial intelligence (AI) tasks dating back from the 1950s (Weaver, 1955). It emphasizes the interaction of language and machine, and how machine can learn human languages with cognitive knowledge (Han, 2022). Before MT, human translation (HT) of written text and documents has always played an important role in science and literature communication between different language speakers, breaking the language barriers. From both MT and HT perspectives, translation quality evaluation (TQE), sometimes also referred to as translation quality assessment (TQA) in variation ¹, is an important task to reflect how well the source text is translated into the target languages (Han et al., 2021b). On the one hand, for low resource language pair scenarios, human translators still play the dominant role in translation production. The translated text and documents can contain unavoidable errors due to personal bias, input efforts, or the training level of the translators. On the other hand, for high resource language pair situation, neural MT (NMT) has achieved remark-

able improvement especially on fluency level, compared to conventional rule-based and statistical phrase-based MT models; however, NMT still has “poisoned cookie” problem struggling to achieve real human parity, for instance, on adequacy level, meaning preservation, and on idiomatic expression translations (Sag et al., 2002; Han et al., 2020b; Johnson et al., 2016; Han et al., 2020a). Translation Service Providers (TSPs) relying on both MT, HT, and human post-editing of MT output (MTPE) carry out translation and editing tasks with the high demand and harsh constraints nowadays. Thus, TQE role in this workflow remains to be critical. However, it is tedious, costly, and time-consuming to check through the entire translated text given the huge amount of data TSPs and customers process.

One obvious solution, to this point, is to extract a subset of the translated text and make a conclusion about the overall translation quality by results of TQE of the sample, which has always been done in real practice. However, one question arises here: *how large the sample size shall be to estimate the overall translation quality of the entire material reliably?* In other words, what is the confidence interval of such evaluation on certain desired confidence level (which is commonly taken as 95%) with the samples we choose to estimate the overall translation quality?

In this work, we carry out such a motivated experimental investigation on confidence evaluation of translation quality evaluation. To take advantage of statistical modeling techniques, we start with the assump-

¹ISO 9001 distinguishes between “assessment” (which relates to the quality management process per se) and evaluation (which is the quality measurement task), so it is incorrect to refer to quality evaluation as “quality assessment” <https://www.iso.org/iso-9001-quality-management.html>.

LH's contribution is done while transferring from ADAPT Centre / Dublin City University to Uni. of Manchester

tion that error distribution is uniform across the entire material, and minimum unit where the error occurs is one sentence (since errors can be between words, in the form of punctuation and conjugation, etc.). This assumption is the best case scenario which potentially ensures the smallest size of the sample, because if it is not correct we immediately arrive to situation when we need to check more, if not entire text. We start from high quality translation assumption that errors are rare (average error density is as low as 0.07 errors per sentence). Then, on the second stage of experiment, by taking post-editing distance (PED) measurement for evaluation, we prove that such assumption can be applied to the situation where errors are much more frequent, with large number of “errors” (edits) per one sentence. We use Bernoulli Statistical Distribution Modeling (BSDM) and Monte Carlo Sampling Analysis (MCSA) to explore the confident level sample size estimation. With this methodology, we expect to reach a conclusion where practical suggestions can be given on confident sample size of translated text/document quality evaluation.

To our knowledge, this is the first study to carry out statistical estimation on confidence intervals for MT quality assessment, while in all practical situations, researchers and practitioners very often take it for granted by randomly choosing a sample size from MT outputs to estimate the system quality.

This paper is organized as below: Section 1 presents the topic of this work, Section 2 covers some related work to our study, Section 3 introduces the statistical modeling, Section 4 presents the Monte Carlo (MC) simulation, Section 5 carries out confidence estimation for post-editing distance metric using MC simulation, and finally Section 6 concludes the work with discussion and future work.

2. Related Work

Quality estimation and measurement have been explored using various strategies for both human and machine translated text. For instance, human evaluation criteria have included intelligibility, fidelity, fluency, adequacy, and comprehension as earlier stage development. These are often carried out based on both source text and candidate translations. Then, editing distance, or post-editing distance are applied to such study by calculating editing steps that are needed to revise the candidate translations towards a correct text sometimes using reference translations. Whenever reference translations are available, automatic reference-based evaluation of the candidate translation quality can be measured via many computational metrics including n-gram based simple word matching as well as some linguistic motivated features such as paraphrasing and Part-of-Speech tagging (Han et al., 2021b). For automatic post-editing distance based metrics, algorithms are designed to calculate the editing steps of insertion, deletion, substitution and/or word reorder-

ing. While automatic metrics have been getting very popular due to their low cost and repeatability, and thus easier to perform comparisons to previously published work, the credibility of such automatic evaluations have always been an issue. For instance, in contrast to the findings from conventional statistical MT researchers (Sanchez-Torron and Koehn, 2016), very recent research work on Neural MT (MT) by (Zouhar et al., 2021; Freitag et al., 2021; Han et al., 2021a) has shown that automatic evaluation metrics such as BLEU do not agree well with human evaluations by professional translators or experts, instead they tend to correlate closely to lower-quality crowd-source based human evaluation. Another disadvantage for automatic evaluation is that we can not get in-depth view of what kind of errors the candidate translations present in the studied context, except for an overall evaluation score, or segment-level scores, not to mention that most metrics do not even allow for clear interpretation of what does the score exactly mean. Regarding this aspect, professional translators can always do a much better job by giving transparent error analysis and categorization on the candidate translations, such as idiomatic expressions (Han et al., 2020a; Han et al., 2020b).

However, another issue arrives at this point, that is how to correctly choose a confident sample from the candidate translation, instead of just take a random sample for granted and try to blindly extrapolate its result to the entire material? Actually this is not a brand new challenge in natural language processing (NLP) field. Having it in mind that randomly chosen samples may contain model bias against a proper evaluation, Prabhu et al. (2019) proposed an uncertainty sampling approach for *text classification* task, and their statistical models can reduce the bias effectively with smaller size of data in comparison to confessional models; Similarly, Haertel et al. (2008) carried out work on how statistical sampling models can help reduce the high cost for Penn Tree-bank annotation while maintaining the higher accuracy; Nadeem et al. (2020) investigated into a systematic comparison of several sampling algorithms used for language generation task, including top-*k*, nucleus and tempered sampling, looking into quality-diversity trade-off. Sampling method was also applied into confident level evaluation of MT. For instance, Koehn (2004) proposed to use bootstrap re-sampling methods to test the significance level of automatic metric BLEU, but using a fixed number of sentences, i.e. 300. Like many other research work, the *chosen number of sentences for evaluation* was never explained or justified with any statistical validation. In contrast, we carry out statistical sampling modeling to estimate the number of sentences that is confident enough to achieve reliable quality evaluation, which means a better representation and generalization of the overall candidate translations in question, with an confidence-cost trade-off.

3. Statistical Modeling

3.1. Study Setup

To carry out a statistical modeling of our research questions under study, i.e. *how to confidently choose a sample size of translation outputs to estimate the overall translation quality, from either HT or MT*, we setup the following initial assumptions:

- Translation errors belong to several independent categories
- Errors of one category are independent from each other.
- Errors of one category occur in text randomly.
- The smallest unit of text where language error occurs is a sentence (in other words, error can be between words, but cannot be between sentences).
- The resulting error distribution is a superposition of distribution of errors of different categories.

Then, if we further project these assumptions into simpler mathematical notes as below, it meets the definition of Bernoulli trial (Brown et al., 2001; Agresti and Coull, 1998):

- Errors of certain type (category) either present in a sentence, or not.
- Errors are independent from each other.
- The probability of errors is the same.

Generally speaking, we cannot always assume that the distribution of errors in the text is uniform, but for the purpose of this exercise it is a reasonable assumption because it is the best possible scenario. The reason is that if errors themselves are not normally distributed across the text, then errors are not distributed uniformly across the material, and the sample size should be much larger to capture clusters of errors, or even 100% review of the material may be required to find such clusters. Therefore, the assumption that errors are uniformly distributed is the best scenario, which allows for the smallest sample size possible.

While this assumption may be questioned, in our experimental evaluations, it proves that the starting value of random seed does not affect the overall model behavior and its solutions. For high quality human translations, these assumptions can be actually very reasonable. Regarding MT outputs and less than premium quality translation scenarios, where many errors per one sentence is not something unusual, we will deal with such more error-prone materials separately in later sections of this paper (Section 4).

3.2. Bernoulli Distribution

In Bernoulli statistical distribution, when the sample size n is significantly smaller than the overall population N , the standard deviation of sample measurement falls into the following formula:

$$\sigma = \sqrt{\frac{p \cdot (1 - p)}{n}}$$

where p is the probability estimation of an event under study. The confidence interval CI , using the Wald interval (Newcombe, 2012)), will be:

$$CI = p \pm \Delta$$

where Δ is the product of standard deviation and factor 1.96 (when confidence level 95% is chosen) (Agresti and Coull, 1998).

$$\Delta = 1.96 \cdot \sigma = 1.96 \cdot \sqrt{\frac{p \cdot (1 - p)}{n}}$$

When the sample size n is comparable to the population size N , e.g. in a smaller translation evaluation project for our study, the standard deviation is calculated as below and the Δ value updates correspondingly:

$$\sigma = \sqrt{\frac{p \cdot (1 - p)}{n} \cdot \frac{N - n}{N - 1}}$$

Let's come back to our study, with this even distribution assumption of each sentence regarding translation errors: having error probability p with value 1 and no error probability $1 - p$ with value 0; each sentence representing a random variable in the modeling. This forms a Binomial distribution $B(n, p)$ where n is the number of sentences.

3.3. Case Studies

We present case studies using both high quality translation text and low quality ones. We first carried out a case study using high quality translation. Statistics from language service providers² shows that the average length of English sentence is 17 words; there are about 250 words on standard page; we therefore can assume that there are 15 English sentences on a standard page. Let's assume that there is a very high quality translation document, where there are no more than one error per page (one error per 15 sentences); then error density $p = 0.07$. And, if we set Δ value as 0.02, the confidence interval falls into 0.07 ± 0.02 , i.e. from 0.05 to 0.09, which is already wide interval. If we use a confidence level 95%, we have the following recommended number of sentences to check as derived from the formula mentioned earlier:

²for instance, <https://logrusglobal.com/>

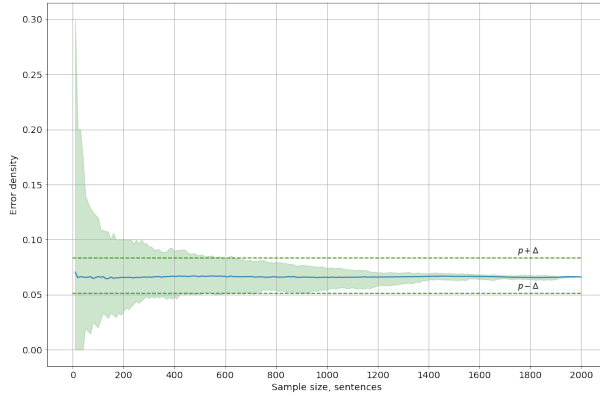


Fig. 1: A 95% confidence level credible interval for sample size from 100 to 2000 sentences, for high quality translation with error density 0.07. Δ is shown for sample size 625 sentences (42 pages).

$$n = \frac{1.96^2 \cdot p \cdot (1 - p)}{\Delta^2}$$

$$n = \frac{1.96^2 \cdot 0.07 \cdot (1 - 0.07)}{0.02^2} = 625$$

which requires us to check 625 sentences, or 42 pages, or 10,000 words (all these numbers are rounded) as demonstrated in Figure 1. If we want to achieve measurement accuracy with confidence interval 0.07 ± 0.01 then we will have to check 2500 sentences, or 166 pages, or 42,000 words. If we check 10 pages the confidence interval will be 0.07 ± 0.04 (from 0.03 to 0.11), and if we check only 1 page: 0.07 ± 0.12 (from 0 to 0.19). If we try to check only one sentence even of high quality translation, we really can't say anything about the quality because the error probability of such measurement will be 0.07 ± 0.5 (from 0 to 0.57) which is too high.

Then let's have a look into examples with lower quality translation text, for instance, an error density of 0.2, i.e., one error for every 5 sentences. The confidence interval of checking only one page will be ± 0.2 : the measurement confidence will be 0.2 ± 0.2 and the actual value can be anywhere from 0 to 0.4 as shown in Figure 2. As we can see, the one page measurement in this situation is not reliable at all.

As we mentioned in Section 3.2, the formulas above only apply for the cases when the sample size is much smaller than the overall population size. If sample size n is comparable to the population size N , the standard deviation of such measurement in the case of high quality translation evaluation will be differently calculated as:

$$\Delta = 1.96 \cdot \sqrt{\frac{p \cdot (1 - p)}{n} \cdot \frac{N - n}{N - 1}}$$

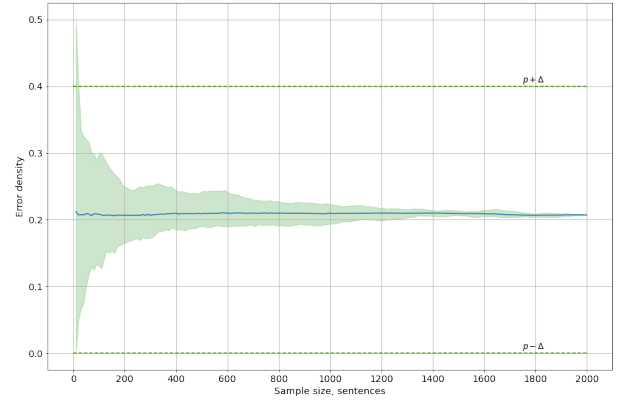


Fig. 2: A 95% confidence level credible interval for sample size from 100 to 2000 sentences, for low quality (machine) translation with error density 0.2. Δ is shown for sample size 15 sentences (1 page).

Practicality of this can be illustrated with the following example: suppose that we have a translation quality evaluation (TQE) scorecard measuring the quality of 1000 words job by doing TQE on 400 words sample checked by reviewer, who found 2 minor errors. Then measured error density is 2 errors per 23.5 sentences, or 0.085 errors per sentence.

The 95% confidence interval for such measurement will be $\pm \Delta$, where:

$$\Delta = 1.96 \cdot \sqrt{\frac{0.085 \cdot (1 - 0.085)}{23.5} \cdot \frac{58.82 - 23.5}{58.82 - 1}} = 0.088$$

Therefore, with such a scorecard measurement we can only say that the measured error density is 0.085 ± 0.088 , i.e. with 95% confidence the actual error density is in the range $[0.0, 0.173]$. Of course, we can check all 1000 words of this job, and then we will be 100% certain about the number of errors in this job (assuming that the reviewer is reliable, but that's another story), but then the measurement confidence interval will relate to the generalized conclusion about the work of this supplier in general, on population size much larger than 1000 words.

In summary, we have the following findings from Bernoulli Distribution modeling for the case studies:

1. If the sample size is much smaller than the entire material, the correct sample size for the measurement does NOT depend on the size of the material, but does depend on the error density in the material and sample size. This case applies to continuous delivery cases, and spot checks on them, or very large projects.
2. If the sample size is comparable with the size of the entire material, the correct sample size for the measurement also depends on the size of the material as well as on the error density in the material and sample size. This case applies to small jobs.

3. The more precisely we want to measure, the more we need to check.

However, in reality, the formulas for Bernoulli distribution will not always apply successfully because that the error distribution in the text is sometimes not uniform. In order to see what is happening with reliability of quality evaluation in more challenging scenarios, in the next section, we conduct numerical Monte-Carlo sampling experiment to examine how our theory comes to practice and then introduce real life complications.

4. Monte Carlo Simulation (MCS)

Monte Carlo Simulation (MCS) was first developed, in the last century when computer machine was still newborn, for a gambling situation when various outcomes can't be easily predicted or determined due to many random variables. MCS assigns some variables from a task that have certain uncertainty with random values, then runs the model to predict a result, which process is subsequently repeated with a great number of times. Finally, estimation will be achieved by averaging the large number of results. These results can also be used to assess statistical properties of error distribution in the samples. The repeated process is also called *sampling* in many applications (Goodfellow et al., 2016) (*Chap. 17*), such as in our translation quality evaluation task under study.

To conduct a numerical experiment we need to take certain parameters of the translation error distribution in a sample of given size.

It is intuitively clear that, the number of errors in one sample will be different each time when the sample is generated simply because errors are distributed randomly throughout the text. This includes the situations of 100% correct judgments of errors and no false positives.

It is less obvious that the parameters of error distribution in the sample are different from parameters of error distribution in the entire material (*aka* "population"). A degenerated case of minimal sample (e.g. one sentence) helps to understand this: if error density in population is 5% and we only analyze one sentence, with 95% probability there will be no errors in that sentence, however it is 100% sure that there are errors in entire material.

The estimation of rare event probabilities are best analyzed with MCS method; sometimes, it is probably the only way to handle such phenomena. As we mentioned, MCS relies on repeated random sampling to obtain numerical results for random variables where models are not available analytically. In our study, it is the case of translation error in the text, with many types of it which interact in a complex way, and the distribution of these errors is not uniform due to the text having a structure of unequal content. The possible reasons for this include that different people have worked on the entire text, and a plethora of other sophisticated reasons.

Our initial experiment will not be dealing with the complexity of many types of errors but examines the simplified model described in earlier section (Section 3). Correspondingly, our numerical MCS experiment to assess parameters of this distribution is described as follows:

1. We take a sample size $N = 2000$ for repeated process.
2. We generate the random distribution of errors in the entire "population" of all the sentences of the material.
3. The number of errors found in these samples represent error distribution of a total number of errors in a sample.
4. We use the large number of sampled data to estimate the entire collection of materials, the error distribution, mean and confidence intervals of such distribution.
5. We take the same error density assumption of 0.07 as in the earlier section (Section 3.3).

After, we will be able to assign different values of these parameters and see how this affects the corresponding results.

Now that we have 2000 samples of text to be evaluated, we can look into the parameters of error distribution in such samples, based on a large dataset. We can calculate the mean of a number of errors in a sample, by creating a list of 2000 elements with one value - the number of errors found in a sample set.

When this is done we have a 2000 elements array of error numbers in 2000 samples, and we can plot the histogram, as shown on Fig. 3. This histogram is a discrete representation of probability density function of distribution of error numbers across the samples. It is not a surprise that the distribution is a Bell curve; regardless of whether errors in the entire population are uniformly distributed or not, the number of errors in the sample is normally distributed, provided the number of samples is sufficiently large, according to the Central Limit Theorem.

The parameters of this distribution can be reconstructed with *norm.fit* function of "scipy.stats"³ programming package, as in Figure 4.

Now, if we take a much smaller (20 times smaller) sample size of 100 sentences we will get a much wider bell, as in Figure 5, which naturally corresponds to wider probability distribution, in Figure 6.

If we run a computational experiment to plot 95% confidence level credible interval for sample sizes from 10 to 2000 sentences, such an experiment will produce results very close to the formulas in previous section (Section 3.3), as shown on Figure 1.

³<https://docs.scipy.org/doc/scipy/reference/stats.html>

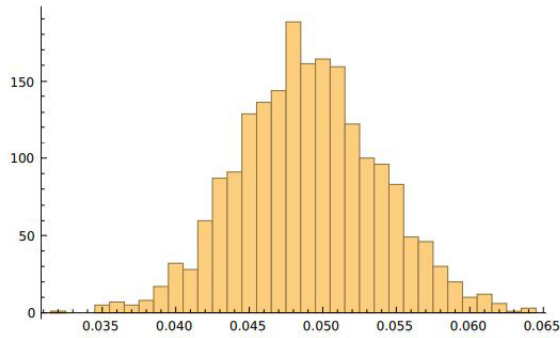


Fig. 3: Distribution histogram of error number in samples, with sample size of 1000 sentences.

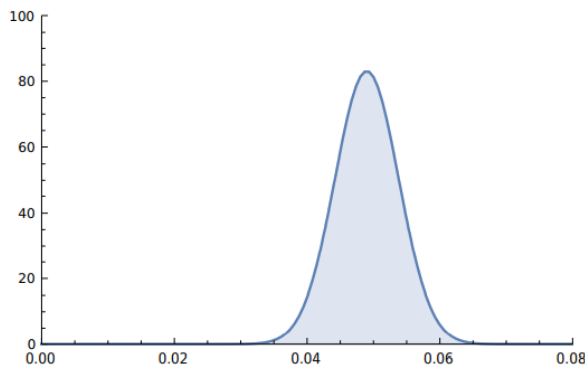


Fig. 4: Reconstructed probability density function of distribution of error number in samples of 1000 sentences.

In the next section, we will demonstrate how MCS can be applied to the confidence estimation on language quality metrics (LQMs), such as post-editing distance.

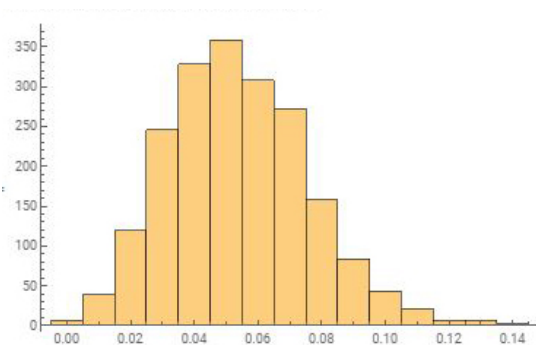


Fig. 5: Distribution histogram of error number in samples, with sample size of 100 sentences.

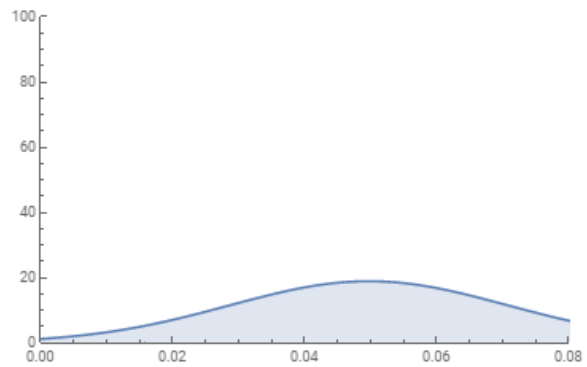


Fig. 6: Reconstructed probability density function of distribution of error number in samples of 100 sentences.

5. MCS for Post-editing Distance (PED) Confidence Estimation

As we mentioned in Section 1, in this section, we investigate how MCS can be applied to simulate translation quality evaluation and confidence estimation task.

Due to the lack of reliable and precise automatic language quality evaluation (LQE) metrics, currently Post-Editing Distance (PED) score remains one of the popular measurement. The PED score is often tracked on segment level, in comparison to document or system level, to examine how good the MT output is in comparison to human edited final translation.

We further conduct analysis of confidence interval (CI, or Δ) for average PED score depending on the sample size.

The absolute PED score is the minimal number of deletion and insertion operation/editing steps from initial candidate translation (*source*) to post-edited translation text, divided by the length of *source* sentence. Because the number editing steps can be larger than the number of words in the *source* sentence, the absolute PED score can be greater than 1⁴.

However, since we design to compare the PED with vector similarity, the absolute PED score needs to be normalized to a [0, 1] range. We propose a normalization function of PED (represented as PED_n) as below:

$$PED_n = 1 - \tanh(c \cdot PED)$$

where c is a parameter defining the value of PED which brings the value of normalized PED to [0, 1], as shown in Figure 7.

Using the MCS methods we introduced in the last section (Section 4), we conduct numerical experiment to estimate confidence interval (Delta) for average normalized measurement PED_n as function of sample size (number of sentences).

⁴the 'source' here indicates the initial candidate translation that is to be edited; pay attention to that it is not the source language text to be translated to the target language.

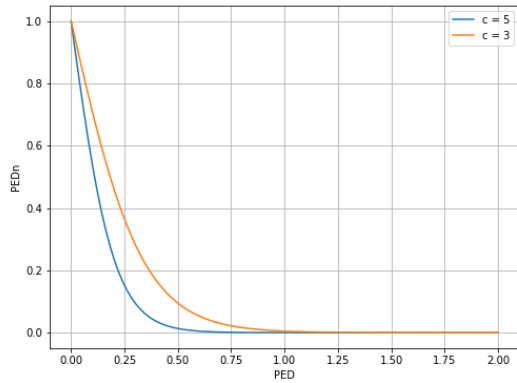


Fig. 7: Relationship between plain PED and normalized PEDn.

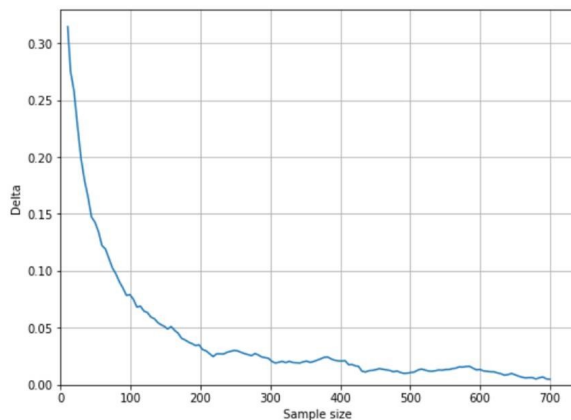


Fig. 8: Delta as function of sample size.

This allows us to understand the number of sentences to calculate average PEDn from samples with confidence interval $\pm \Delta$ value.

First we simulate Delta (half of confidence interval) as function of the sample size, as in Figure 8. Then the confidence interval is shown in Figure 9 which graphic indicates that confidence interval starts going out of control when the sample size is less than 200 sentences (approximately 3500 words).

This means that reliable conclusions can not be reached about the quality of MT proposals based on post-editing sample sizes of 2000 words or less. A sample size above 4000 words is recommended to get reasonable confidence interval in a better scenario, for purely statistical reasons, set aside subjective difference of opinions of different editors, content irregularities, and other possible factors.

6. Discussions and Conclusions

In this article, we proposed the research question on confident sample size of translation documents to es-

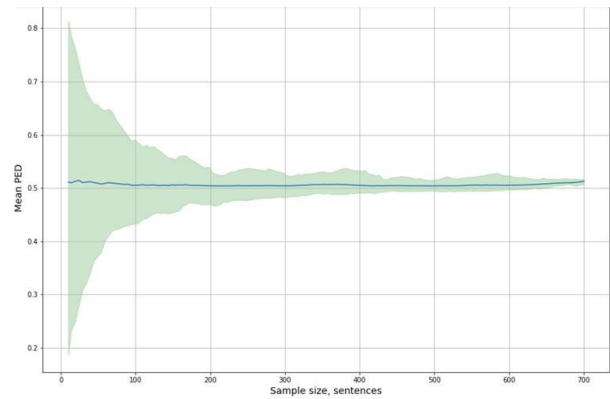


Fig. 9: A 95% confidence level credible interval for mean PEDn.

timate the overall material quality, which is a crucial question for both academic research and industrial applications, such as for clients and language service providers.

We started the experimental investigation of translation quality evaluation (TQE) by assuming that the errors in translated text are evenly distributed, with errors being rare (7 errors per 100 sentences, on average). This assumption is placed as a random seed of our statistical *Bernoulli* modeling which does not affect the overall model behavior. To simulate the practical situation where the errors can come from different translators and different types, and span into a different weight across the translated text and documents, we applied the *Monte Carlo* Simulation analysis, using a sample size of 2000 sentences and 95% confidence level. We also applied MCS into confidence estimation of post-editing distance measurement which is currently a widely adopted evaluation metric for translation assessment, and gained very valuable findings from empirical investigations regarding practical situations when translation quality evaluation (TQE) is deployed.

Our statistical modeling using MCS suggests that a sample size of less than 200 sentences can not reflect the overall material quality in a confident enough level on translation quality evaluation task. Using MCS, we also reduced the suggested sample size from 10k words (around 625 sentences, from Bernoulli statistics in Section 3) to 4k (Section 5) for reliable estimation of overall translation quality.

Furthermore, we suggest that, ideally, a reliability level of analytic sample quality measurement can be added to every analytic TQE scorecard in the form of confidence interval at certain confidence level as one important indicator of the level of certainty of measurement results. In the future work, we plan to compare different sampling methods, as well as apply the confidence estimation model into broader TQE metrics.

7. Acknowledgments

We thank the anonymous reviewers for very valuable comments and suggestions on our work. Lifeng Han’s contribution was partially funded by The University of Manchester and ADAPT Research Centre: the ADAPT Centre for Digital Content Technology is funded under the Science Foundation Ireland (SFI) Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

8. Bibliographical References

- Agresti, A. and Coull, B. A. (1998). Approximate is better than “exact” for interval estimation of binomial proportions. *The American Statistician*, 52(2):119–126.
- Brown, L. D., Cai, T. T., and DasGupta, A. (2001). Interval Estimation for a Binomial Proportion. *Statistical Science*, 16(2):101 – 133.
- Freitag, M., Foster, G., Grangier, D., Ratnakar, V., Tan, Q., and Macherey, W. (2021). Experts, Errors, and Context: A Large-Scale Study of Human Evaluation for Machine Translation. *arXiv e-prints*, page arXiv:2104.14478, April.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Haertel, R., Ringger, E., Seppi, K., Carroll, J., and McClanahan, P. (2008). Assessing the costs of sampling methods in active learning for annotation. In *Proceedings of ACL-08: HLT, Short Papers*, pages 65–68, Columbus, Ohio, June. Association for Computational Linguistics.
- Han, L., Jones, G., and Smeaton, A. (2020a). AlphaMWE: Construction of multilingual parallel corpora with MWE annotations. In *Proceedings of the Joint Workshop on Multiword Expressions and Electronic Lexicons*, pages 44–57, online, December. Association for Computational Linguistics.
- Han, L., Jones, G., and Smeaton, A. (2020b). MultiMWE: Building a multi-lingual multi-word expression (MWE) parallel corpora. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2970–2979, Marseille, France, May. European Language Resources Association.
- Han, L., Jones, G., Smeaton, A., and Bolzoni, P. (2021a). Chinese character decomposition for neural MT with multi-word expressions. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 336–344, Reykjavik, Iceland (Online), May 31–2 June. Linköping University Electronic Press, Sweden.
- Han, L., Smeaton, A., and Jones, G. (2021b). Translation quality assessment: A brief survey on manual and automatic methods. In *Proceedings for the First Workshop on Modelling Translation: Translatology in the Digital Age*, pages 15–33, online, May. Association for Computational Linguistics.
- Han, L. (2022). *An Investigation into Multi-Word Expressions in Machine Translation*. PhD thesis, Dublin City University. <https://doras.dcu.ie/26559/>.
- Johnson, M., Schuster, M., Le, Q. V., Krikun, M., Wu, Y., Chen, Z., Thorat, N., Viégas, F. B., Wattenberg, M., Corrado, G., Hughes, M., and Dean, J. (2016). Google’s multilingual neural machine translation system: Enabling zero-shot translation. *CoRR*, abs/1611.04558.
- Koehn, P. (2004). Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain, July. Association for Computational Linguistics.
- Nadeem, M., He, T., Cho, K., and Glass, J. (2020). A systematic characterization of sampling algorithms for open-ended language generation. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 334–346, Suzhou, China, December. Association for Computational Linguistics.
- Newcombe, R. G. (2012). *Confidence Intervals for Proportions and Related Measures of Effect Size*. CRC Press. <https://doi.org/10.1201/b12670>.
- Prabhu, A., Dognin, C., and Singh, M. (2019). Sampling bias in deep active classification: An empirical study. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4058–4068, Hong Kong, China, November. Association for Computational Linguistics.
- Sag, I. A., Baldwin, T., Bond, F., Copestake, A., and Flickinger, D. (2002). Multiword expressions: A pain in the neck for nlp. In Alexander Gelbukh, editor, *Computational Linguistics and Intelligent Text Processing*, pages 1–15, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Sanchez-Torron, M. and Koehn, P. (2016). Machine translation quality and post-editor productivity. In *Conferences of the Association for Machine Translation in the Americas: MT Researchers’ Track*, pages 16–26, Austin, TX, USA, October 28 - November 1. The Association for Machine Translation in the Americas.
- Weaver, W. (1955). Translation. *Machine Translation of Languages: Fourteen Essays*.
- Zouhar, V., Tamchyna, A., Popel, M., and Bojar, O. (2021). Neural machine translation quality and post-editing performance. In *EMNLP-2021*.