

SurreyAI 2023 Submission for the Quality Estimation Shared Task

Archchana Sindhujan^{*1}, Diptesh Kanojia¹, Constantin Orasan² and Tharindu Ranasinghe³

¹ Institute for People-Centred AI, University Of Surrey , UK

²Centre for Translation Studies, University of Surrey, UK

³ School of Computer Science and Digital Technologies, Aston University, UK

{a.sindhujan,d.kanojia,c.orasan}@surrey.ac.uk,

t.ranasinghe@aston.ac.uk

Abstract

Quality Estimation (QE) systems are important in situations where it is necessary to assess the quality of translations, but there is no reference available. This paper describes the approach adopted by the SurreyAI team for addressing the Sentence-Level Direct Assessment shared task in WMT23. The proposed approach builds upon the TransQuest framework, exploring various autoencoder pre-trained language models within the MonoTransQuest architecture using single and ensemble settings. The autoencoder pre-trained language models employed in the proposed systems are XLMV, InfoXLM-large, and XLMR-large. The evaluation utilizes Spearman and Pearson correlation coefficients, assessing the relationship between machine-predicted quality scores and human judgments for 5 language pairs (English-Gujarati, English-Hindi, English-Marathi, English-Tamil and English-Telugu). The MonoTQ-InfoXLM-large approach emerges as a robust strategy, surpassing all other individual models proposed in this study by significantly improving over the baseline for the majority of the language pairs.

1 Introduction

The primary objective of quality estimation (QE) systems is to assess the quality of a translation without relying on a reference translation. This makes QE valuable within translation processes, as it enables the determination of whether an automatically generated translation is sufficiently accurate for a specific purpose. This aids in deciding whether the translation can be used as is, requires human intervention for full translation, or necessitates post-editing by a human translator (Kepler et al., 2019b). Quality estimation can be conducted across various levels: word/phrase level, sentence level and document level. This paper considers only the sentence-level QE and presents our participation in the WMT23 Sentence-level direct assessment (DA) shared task. In

the context of this task, participating systems are required to predict the DA score for a given (source, target) pair. This score serves as a measure of the translation quality.

Building upon the ideas presented in TransQuest by Ranasinghe et al. (2020b), our investigation explores the use of various pre-trained models within the MonoTransQuest architecture for the sentence-level quality estimation shared task. The architecture employs autoencoder pre-trained language models to fine-tune the QE data to predict a score which indicates the quality of translation. Using the MonoTransQuest architecture as the base we employ the pre-trained transformers separately to implement the systems MonoTQ-XLMV, MonoTQ-InfoXLM-large and MonoTQ-XLMR-large. In addition, we propose ensembleTQ which combines the output of MonoTransQuest when using different pre-trained models. All the proposed systems achieve a significantly higher Spearman correlation score compared to the baseline.

The paper is structured as follows. Section 2 briefly presents related work on quality estimation. Section 3 provides a concise overview of the dataset used in the sentence-level QE shared task. Moving on to Section 4, we introduce the autoencoder pre-trained language models and proposed systems and detail the training methodology. Section 5 is dedicated to the evaluation and Section 6 comprises the result and discussion. The paper concludes by summarizing findings, highlighting conclusions, and suggesting potential avenues for future research in the final section.

2 Related work

Quality estimation in machine translation has evolved significantly throughout the years. Initially, it relied on feature engineering and conventional machine learning techniques like SVM and basic

neural networks (Specia et al., 2015; Scarton and Specia, 2014). However, Neural Networks has since become central to quality estimation, where there is no more need of feature engineering, and the models can be trained directly on the data (Kepler et al., 2019b,a; Specia et al., 2018). Recently, Transformer-based architectures have arisen as robust solutions for machine translation and quality estimation. Notably, there are two widely recognized frameworks that leverage this transformative approach for QE tasks: TransQuest (Ranasinghe et al., 2020a) and CometKIWI (Rei et al., 2022).

Ensemble methods have also been explored extensively in Quality Estimation tasks (Bao et al., 2022; Geng et al., 2022; Kepler et al., 2019b; Ranasinghe et al., 2020b; Rei et al., 2022). The ensemble approach from Lim and Park (2022) using K-folds consistently outperformed the standard method, underscoring the prevalent belief that ensemble strategies enhance performance outcomes. However, some of the research studies (Ranasinghe et al., 2020b; Bao et al., 2022; Rei et al., 2022) show that combining multi-lingual models through ensembling yields better results than the traditional k-fold ensemble technique. Geng et al. (2022) suggest an alternative ensemble method that merges the results from models trained using various sentence-level metrics.

Our study delves into the performance of cutting-edge pre-trained transformer-based approaches when applied to sentence-level Quality Estimation tasks.

3 Dataset

We focus on Sentence-Level Direct Assessment tasks which comprise datasets for 5 language pairs which has English on the source side and Indian languages on the target side: English-Gujarati (En-Gu), English-Hindi (En-Hi), English-Marathi (En-Mr), English-Tamil (En-Ta) and English-Telugu (En-Te). Among these language pairs, En-Hi language pair is considered mid-resourced and all the other language pairs are low-resourced. Each language pair includes around 7,000 sentence pairs in the training set, as well as around 1,000 sentence pairs in both the development and testing sets. Each translation was evaluated by three professional translators who assigned a score between 0 and 100. These Direct Assessment (DA) scores were normalized using the z-score. The final score for

the sentence-level task requires predicting the mean DA z-scores for the test sentence pairs. More details on this can be found in Zerva et al. (2022).

4 Methodology

This section outlines the approach taken to formulate our quality estimation techniques. We begin by detailing the autoencoder pre-trained language models employed in our architecture. Then we explain the architecture and the strategy employed to train these network architectures in detail.

4.1 Pre-trained models for fine-tuning

1. XLMR-large

XLM-Roberta (Conneau et al., 2020) is a pre-trained transformer-based language model which is a part of the Cross-lingual Language Model (XLM). This model employs large-scale cross-lingual pre-training to capture contextual information and representations across 100 languages. The model is trained on 2.5TB of filtered CommonCrawl data from multiple languages, allowing it to effectively learn cross-lingual and language-specific patterns. The XLM-R architecture takes sequences as input, with a maximum token limit of 512, and generates contextualized embeddings for each token, enabling it to perform well on various natural language processing tasks across different languages (Ranasinghe et al., 2020b, 2021).

2. XLMV

XLMV is a multilingual language model with a one million token vocabulary trained on 2.5TB of data from Common Crawl (same as XLM-R) (Liang et al., 2023). In the context of large multilingual language models, a common practice involves employing a single vocabulary shared across a diverse set of languages. Even with the expansion in model complexity, including parameter count and depth, the vocabulary size has remained relatively static. This constraint in vocabulary hampers the potential of multilingual models such as XLM-R to capture nuanced representations effectively (Wang et al., 2019).

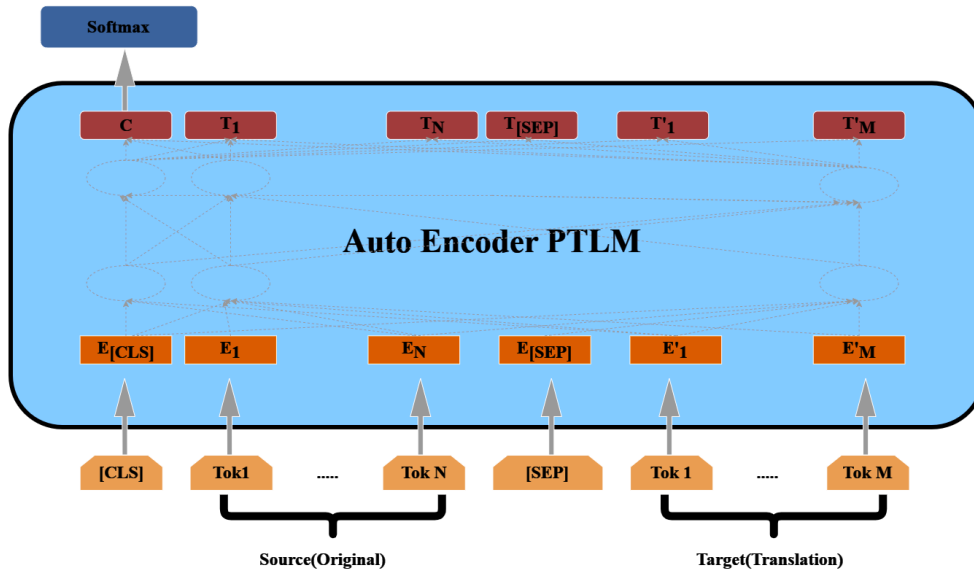


Figure 1: Architecture diagram of the proposed approaches

XMLV introduced an innovative strategy addressing this issue by achieving scalability to extensive multilingual vocabularies. XMLV involves prioritizing vocabulary allocation based on language-specific lexical overlap, ensuring sufficient coverage for each language. The outcome is tokenizations that hold enhanced semantic significance and are generally more concise compared to those generated by XLM-R.

3. InfoXLM-large

InfoXLM-large (Chi et al., 2021), is an information-theoretic framework for cross-lingual language model pre-training. It extends the XLM-R architecture by formulating cross-lingual pre-training to maximize mutual information between multilingual texts at different granularities. This approach enhances the model’s capability to learn effective cross-lingual representations by capturing shared information across languages. InfoXLM-large introduces a novel pre-training task based on contrastive learning, treating bilingual sentence pairs as views of the same meaning. By jointly training on monolingual and parallel corpora, the model improves the transferability of its representations for various downstream cross-lingual tasks (Rei et al., 2022; Bao et al., 2022).

4.2 Architecture

The proposed architecture of MonoTransQuest employs a pre-trained language model as shown in Figure 1. The MonoTransQuest architecture in TransQuest (Ranasinghe et al., 2020b) considers only the XLMR transformer model. In our proposed system, we train multiple multilingual QE models by fine-tuning autoencoder pre-trained language models (PTLMs) and report mean z-scores. The PTLMs are namely XMLV, InfoXLM-large, and XLMR-large which we have explained in section 4.1. The model’s input consists of the original sentence (source) and its translation (target) concatenated, with a [SEP] token. This token marks the separation of the original sentence and the translated sentence. The pre-trained autoencoder accepts input sequences with a token limit of 512 and produces a sequence representation as output. The initial token of the sequence is [CLS] token, encompassing a distinctive embedding to signify the entire sequence. Subsequently, embeddings are assigned to each word in the sequence. Ranasinghe et al. (2020b) highlights the superiority of the CLS-strategy over the MEAN-strategy (calculating the mean of all output vectors corresponding to the input words) and MAX-strategy (determining the maximum value across the output vectors of input words) for pooling within the MonoTransQuest framework. We have used the CLS-strategy (using the output of the [CLS] token) to extract the output from the transformer model. Consequently, we employed

Method	En-Gu		En-Hi		En-Mr		En-Ta		En-Te	
	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r
I Baseline	0.337	0.307	0.281	0.245	0.392	0.427	0.507	0.402	0.193	0.153
II MonoTQ-XLMV	0.673	0.536	0.572	0.687	0.642	0.425	0.670	0.559	0.464	0.642
III MonoTQ-InfoXLM-large	0.713	0.656	0.624	0.726	0.470	0.030	0.726	0.662	0.462	0.719
IV MonoTQ-XLMR-large	0.438	0.299	0.440	0.430	0.395	-0.117	0.482	0.454	0.345	0.211
V ensembleTQ	0.649	0.700	0.551	0.668	0.596	0.668	0.674	0.710	0.349	0.376

Table 1: Spearman (ρ) and Pearson (r) correlation between the proposed approach predictions and human DA judgments. The best Spearman score obtained for each language pair (any method) is marked in bold. Rows II, III, and IV indicate the single-configuration settings of MonoTransQuest architecture with different pre-trained transformer models as explained in Section 5.1, and ensembleTQ in row V is explained in Section 5.2. The baseline results are in Row I.

the [CLS] token’s embedding as input for a softmax layer. The softmax layer predicts the translation’s quality score. The mean-squared-error loss function was used as the objective function for training.

4.3 Training and Implementation Details

We started the training with MonoTQ-XLMV which incorporates the XLMV-base model with MonoTransQuest for all 5 language pairs. We had the batch size as 8. We have used Adam Optimizer (Kingma and Ba, 2014) with a learning rate of $2e-5$. The model is trained using 3 epochs. The training process exclusively utilized the training data. Early stopping was enforced if the evaluation loss failed to show improvements over ten consecutive evaluation rounds. We continued the training with the same set of configurations for MonoTQ-InfoXLM-large and MonoTQ-XLMR-large separately. MonoTQ-XLMV and MonoTQ-InfoXLM-large required twice the training time compared to MonoTQ-XLMR-large which required approximately 40 minutes of training on a GPU with 48GB of memory.

The proposed systems are built upon the most up-to-date version of TransQuest¹ framework and executed using Python 3.9 and PyTorch 2.0.1. The integration of pre-trained encoders (XLMV², XLMR-large³ and InfoXLM-large⁴) into the MonoTransQuest architecture was facilitated

¹<https://github.com/tharindudr/transQuest>

²<https://huggingface.co/facebook/xlm-v-base>

³<https://huggingface.co/xlm-roberta-large>

⁴<https://huggingface.co/microsoft/infoclm-large>

through the application of HuggingFace’s Transformers library.

5 Evaluation

In this section, we outline the evaluation outcomes of our models. We assess the performance of the proposed models under two circumstances: single model configuration and ensembleTQ.

The primary evaluation criterion employed was Spearman’s rank correlation coefficient (Sedgwick, 2014), which is a statistical measure used to evaluate the strength and direction of association between two variables. Also, we have calculated the Pearson correlation coefficient (Cohen et al., 2009) as a secondary metric for the evaluation. In the context of Quality Estimation (QE) for machine translation, it is used to evaluate the correlation between the machine-predicted quality scores and the gold standard labels provided by human annotators in the test dataset. Spearman’s rank correlation coefficient assesses the monotonic relationship between the two variables, unlike the Pearson correlation (Cohen et al., 2009), which measures the linear relationship between two variables. It is calculated by first ranking the values of both variables in ascending or descending order and then computing the Pearson correlation coefficient between the two sets of ranks. Spearman’s rank correlation coefficient is often preferred because it is less sensitive to outliers and non-linear relationships between the predicted scores and human scores.

Team	En-Gu		En-Hi		En-Mr		En-Ta		En-Te	
	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r
1 Unbabel-IST	0.714	0.745	0.598	0.667	0.704	0.735	0.739	0.733	0.388	0.362
2 IOL Research	0.695	0.742	0.6	0.667	0.505	0.372	0.74	0.742	0.376	0.344
3 HW-TSC	0.691	0.714	0.644	0.72	0.692	0.718	0.775	0.778	0.394	0.35
4 MMT	0.54	0.581	0.494	0.57	0.65	0.663	0.547	0.531	0.337	0.281
5 Baseline	0.337	0.307	0.281	0.245	0.392	0.427	0.507	0.402	0.193	0.153
6 SurreyAI-ensembleTQ	0.649	0.700	0.551	0.668	0.596	0.668	0.674	0.710	0.349	0.376

Table 2: Spearman (ρ) and Pearson (r) correlation between the predictions from the participated systems in WMT23 sentence-level QE shared task and human DA judgments. The best Spearman and Pearson score obtained for each language pair is marked in bold. Even though we have experimented with the single model configurations, we only submitted our ensembled approach (SurreyAI-ensembleTQ) for the shared task competition.

5.1 Single model configurations

Initially, our evaluation focused on the single model configurations of the proposed framework. This involved training a quality estimation model using a single autoencoder pre-trained language model on the training data for each language pair separately. Subsequently, we assessed each model’s performance (MonoTQ-XLMV, MonoTQ-InfoXLM-large, MonoTQ-XLMR-large) on the corresponding test set for each language pair. The outcomes of this evaluation for the single model configuration are presented in Table 1.

5.2 EnsembleTQ

Recently, ensemble techniques have demonstrated their efficacy in enhancing transformer-based models’ performance (Xu et al., 2020). Following this approach, we employed an ensemble strategy to experiment further to see whether it enhance the performance. For every input within the test set, we aggregate the output scores from various distinct pre-trained models integrated into the MonoTransQuest architecture. Subsequently, we calculate the average of the cumulative score, divided by the number of pre-trained models, resulting in the ensembleTQ score. Finally, we compute the Spearman and Pearson correlation scores for the ensembleTQ score, providing a comprehensive evaluation of our ensemble approach.

6 Result and Discussion

The research is divided into two distinct settings, as outlined in Sections 5.1 and 5.2. The primary evaluation metric employed in this study is the Spearman correlation coefficient.

As shown in Table 1, is notable that the baseline model does not surpass our proposed approaches in terms of Spearman correlation scores in most cases. This outcome underscores the specific strengths and limitations associated with different model architectures. Complementing the Spearman correlation analysis, the examination of Pearson correlation scores further enriches the assessment. The MonoTQ-InfoXLM-large model consistently exhibits superior Pearson correlation scores across a majority of the language pairs, accentuating its robust performance characteristics.

From our experiment results, as shown in Table 1, it’s notable that the single-model configuration of MonoTQ-InfoXLM-large and MonoTQ-XLMV outperform ensemble-TQ for the majority of the language pairs. Observing the results outlined in both Table 1 and Table 2, it becomes evident that MonoTQ-InfoXLM-large and MonoTQ-XLMV not only outperform other systems among our own proposed approaches, they also exhibit a competitive performance with the best-performing system in the WMT23 sentence-level shared-task. MonoTQ-InfoXLM-large shows a very close Spearman correlation score with the winning system of the WMT23 sentence-level task for the En-Gu, En-Hi and En-Ta language pairs. Also, MonoTQ-XLMV shows the highest Spearman

No.	Name	DiskFootPrint (Bytes)
1	Unbabel-IST	42,868,104,221
2	IOL Research	2,357,242,105
3	HW-TSC	27,730,527,504
4	MMT	2,448,132,038
5	SurreyAI-ensembleTQ	7,945,689,496
6	SurreyAI-MonoTQ-XLMV	3,221,225,472
7	SurreyAI-MonoTQ-InfoXLM-large	2,362,232,012
8	SurreyAI-MonoTQ-XLMR-large	2,254,857,830

Table 3: Rows 1-5 display the disk footprint of ensemble model submissions related to the sentence-level task for WMT23. Meanwhile, Rows 6-8 present the disk footprint of our TQ models with single model configuration.

correlation score for the En-Te language pair. This observation raises the question that do the practice of ensembling always guarantees performance enhancement. Table 3 presents the memory requirements of both ensemble approaches and single-model configurations. Interestingly, in most cases ensemble models demand significantly more memory space than single-model setups, despite only offering a marginal boost in performance. This observation prompts us to reconsider the efficiency of employing ensemble models.

The conducted experiments across mid-resourced and low-resourced language pairs unravel intricate performance dynamics among various models.

7 Conclusion

This paper comprehensively evaluates the proposed architecture within the context of sentence-level direct quality assessment, employing diverse encoder-based pre-trained models. Our investigation notably highlights the enhanced performance attributed to the MonoTQ-InfoXLM-large, which surpasses the other configuration approaches, namely MonoTQ-XLMV, ensembleTQ strategy and

MonoTQ-XLMR-large. While our outcomes in the WMT23 sentence-level Direct Assessment task did not attain peak performance, they nevertheless exhibited a marked improvement over the baseline and showed notable performance scores close to the winning systems.

Looking ahead, our research trajectory anticipates a continued exploration of quality estimation employing large language models. This involves further experimentation encompassing a broader spectrum of low-resourced language pairs. These forthcoming endeavours aspire to deepen our insights into the intricacies of direct quality assessment and contribute to advancing the frontiers of natural language processing. Also, we are focused on continuing the experimentation of pre-trained language models incorporated into different QE frameworks.

References

- Keqin Bao, Yu Wan, Dayiheng Liu, Baosong Yang, Wenqiang Lei, Xiangnan He, Derek F. Wong, and Jun Xie. 2022. [Alibaba-translate China’s submission for WMT 2022 quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 597–605, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Xiang Geng, Yu Zhang, Shujian Huang, Shimin Tao, Hao Yang, and Jiajun Chen. 2022. [NJUNLP’s participation for the WMT2022 quality estimation](#)

- shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 615–620, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, António Góis, M. Amin Farajian, António V. Lopes, and André F. T. Martins. 2019a. [Unbabel’s participation in the WMT19 translation quality estimation shared task](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 78–84, Florence, Italy. Association for Computational Linguistics.
- Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, and André F. T. Martins. 2019b. [OpenKiwi: An open source framework for quality estimation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 117–122, Florence, Italy. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Davis Liang, Hila Gonen, Yuning Mao, Rui Hou, Naman Goyal, Marjan Ghazvininejad, Luke Zettlemoyer, and Madian Khabsa. 2023. [Xlm-v: Overcoming the vocabulary bottleneck in multilingual masked language models](#). *ArXiv*, abs/2301.10472.
- Seunghyun Lim and Jeonghyeok Park. 2022. [Papago’s submission to the WMT22 quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 627–633, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020a. [TransQuest at WMT2020: Sentence-level direct assessment](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1049–1055, Online. Association for Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020b. [TransQuest: Translation quality estimation with cross-lingual transformers](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5070–5081, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2021. [An exploratory analysis of multilingual word-level quality estimation with cross-lingual transformers](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 434–440, Online. Association for Computational Linguistics.
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Carolina Scarton and Lucia Specia. 2014. [Document-level translation quality estimation: exploring discourse and pseudo-references](#). In *Proceedings of the 17th Annual Conference of the European Association for Machine Translation*, pages 101–108, Dubrovnik, Croatia. European Association for Machine Translation.
- Philip Sedgwick. 2014. Spearman’s rank correlation coefficient. *Bmj*, 349.
- Lucia Specia, Frédéric Blain, Varvara Logacheva, Ramón F. Astudillo, and André F. T. Martins. 2018. [Findings of the WMT 2018 shared task on quality estimation](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 689–709, Belgium, Brussels. Association for Computational Linguistics.
- Lucia Specia, Gustavo Paetzold, and Carolina Scarton. 2015. [Multi-level translation quality prediction with QuEst++](#). In *Proceedings of ACL-IJCNLP 2015 System Demonstrations*, pages 115–120, Beijing, China. Association for Computational Linguistics and The Asian Federation of Natural Language Processing.
- Hai Wang, Dian Yu, Kai Sun, Jianshu Chen, and Dong Yu. 2019. [Improving pre-trained multilingual model with vocabulary expansion](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 316–327, Hong Kong, China. Association for Computational Linguistics.
- Yige Xu, Xipeng Qiu, Ligao Zhou, and Xuanjing Huang. 2020. [Improving bert fine-tuning via self-ensemble and self-distillation](#). *arXiv preprint arXiv:2002.10345*.
- Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Steffen Eger, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.