

An Exploratory Analysis of Multilingual Word-Level Quality Estimation with Cross-Lingual Transformers

Tharindu Ranasinghe[◇], Constantin Orăsan[♡] and Ruslan Mitkov[◇]

[◇]Research Group in Computational Linguistics, University of Wolverhampton, UK

[♡]Centre for Translation Studies, University of Surrey, UK

{t.d.ranasinghehettiarachchige, r.mitkov}@wlv.ac.uk

c.orasan@surrey.ac.uk

Abstract

Most studies on word-level Quality Estimation (QE) of machine translation focus on language-specific models. The obvious disadvantages of these approaches are the need for labelled data for each language pair and the high cost required to maintain several language-specific models. To overcome these problems, we explore different approaches to multilingual, word-level QE. We show that multilingual QE models perform on par with the current language-specific models. In the cases of zero-shot and few-shot QE, we demonstrate that it is possible to accurately predict word-level quality for any given new language pair from models trained on other language pairs. Our findings suggest that the word-level QE models based on powerful pre-trained transformers that we propose in this paper generalise well across languages, making them more useful in real-world scenarios.

1 Introduction

Quality Estimation (QE) is the task of assessing the quality of a translation without having access to a reference translation (Specia et al., 2009). Translation quality can be estimated at different levels of granularity: word, sentence and document level (Ive et al., 2018). So far the most popular task has been sentence-level QE (Specia et al., 2020), in which QE models provide a score for each pair of source and target sentences. A more challenging task, which is currently receiving a lot of attention from the research community, is word-level quality estimation. This task provides more fine-grained information about the quality of a translation, indicating which words from the source have been incorrectly translated in the target, and whether the words inserted between these words are correct (good vs bad gaps). This information can be useful for post-editors by indicating the parts of a sentence on which they have to focus more.

Word-level QE is generally framed as a supervised ML problem (Kepler et al., 2019; Lee, 2020) trained on data in which the correctness of translation is labelled at word-level (i.e. good, bad, gap). The training data publicly available to build word-level QE models is limited to very few language pairs, which makes it difficult to build QE models for many languages. From an application perspective, even for the languages with resources, it is difficult to maintain separate QE models for each language since the state-of-the-art neural QE models are large in size (Ranasinghe et al., 2020b).

In our paper, we address this problem by developing multilingual word-level QE models which perform competitively in different domains, MT types and language pairs. In addition, for the first time, we propose word-level QE as a zero-shot cross-lingual transfer task, enabling new avenues of research in which multilingual models can be trained once and then serve a multitude of languages and domains. The main contributions of this paper are the following:

- i We introduce a simple architecture to perform word-level quality estimation that predicts the quality of the words in the source sentence, target sentence and the gaps in the target sentence.
- ii We explore multilingual, word-level quality estimation with the proposed architecture. We show that multilingual models are competitive with bilingual models.
- iii We inspect few-shot and zero-shot word-level quality estimation with the bilingual and multilingual models. We report how the source-target direction, domain and MT type affect the predictions for a new language pair.
- iv We release the code and the pre-trained models as part of an open-source framework¹.

¹Documentation is available on <http://tharindu.co.uk/TransQuest/>

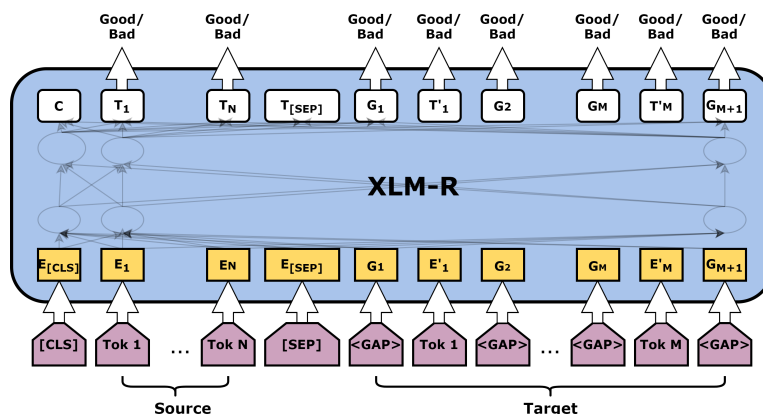


Figure 1: Model Architecture

2 Related Work

Quality Estimation Early approaches in word-level QE were based on features fed into a traditional machine learning algorithm. Systems like QuEst++ (Specia et al., 2015) and MARMOT (Logacheva et al., 2016) were based on features used with Conditional Random Fields to perform word-level QE. With deep learning models becoming popular, the next generation of word-level QE algorithms were based on bilingual word embeddings fed into deep neural networks. Such approaches can be found in OpenKiwi (Kepler et al., 2019). However, the current state of the art in word-level QE is based on transformers like BERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) where a simple linear layer is added on top of the transformer model to obtain the predictions (Lee, 2020). All of these approaches consider quality estimation as a language-specific task and build a different model for each language pair. This approach has many drawbacks in real-world applications, some of which are discussed in Section 1.

Multilinguality Multilinguality allows training a single model to perform a task from and/or to multiple languages. Even though this has been applied to many tasks (Ranasinghe and Zampieri, 2020, 2021) including NMT (Nguyen and Chiang, 2017; Aharoni et al., 2019), multilingual approaches have been rarely used in QE (Sun et al., 2020). Shah and Specia (2016) explore QE models for more than one language where they use multitask learning with annotators or languages as multiple tasks. They show that multilingual models led to marginal improvements over bilingual ones with a traditional black-box, feature-based approach. In a recent

study, Ranasinghe et al. (2020b) show that multilingual QE models based on transformers trained on high-resource languages can be used for zero-shot, sentence-level QE in low-resource languages. In a similar architecture, but with multi-task learning, Sun et al. (2020) report that multilingual QE models outperform bilingual models, particularly in less balanced quality label distributions and low-resource settings. However, these two papers are focused on sentence-level QE and to the best of our knowledge, no prior work has been done on multilingual, word-level QE models.

3 Architecture

Our architecture relies on the XLM-R transformer model (Conneau et al., 2020) to derive the representations of the input sentences. XLM-R has been trained on a large-scale multilingual dataset in 104 languages, totalling 2.5TB, extracted from the CommonCrawl datasets. It is trained using only RoBERTa’s (Liu et al., 2019) masked language modelling (MLM) objective. XLM-R was used by the winning systems in the recent WMT 2020 shared task on sentence-level QE (Ranasinghe et al., 2020a; Lee, 2020; Fomicheva et al., 2020). This motivated us to use a similar approach for word-level QE.

Our architecture adds a new token to the XLM-R tokenizer called <GAP> which is inserted between the words in the target. We then concatenate the source and the target with a [SEP] token and we feed them into XLM-R. A simple linear layer is added on top of word and <GAP> embeddings to predict whether it is “Good” or “Bad” as shown in Figure 1. The training configurations and the system specifications are presented in the supplementary material.

Language Pair	Source	MT System	Competition	Train Size
De-En	Pharmaceutical	Phrase-based SMT	WMT 2018	25,963
En-Cs	IT	Phrase-based SMT	WMT 2018	40,254
En-De	Wiki	fairseq-based NMT	WMT 2020	7,000
En-De	IT	fairseq-based NMT	WMT 2019	13,442
En-De	IT	Phrase-based SMT	WMT 2018	26,273
En-Ru	IT	Online NMT	WMT 2019	15,089
En-Lv	Pharmaceutical	Attention-based NMT	WMT 2018	12,936
En-Lv	Pharmaceutical	Phrase-based SMT	WMT 2018	11,251
En-Zh	Wiki	fairseq-based NMT	WMT 2020	7,000

Table 1: Information about the language pairs used to predict word-level quality. The **Language Pair** column lists the language pairs we used in ISO 639-1 codes. **Source** stands for the domain of the sentence and **MT System** is the Machine Translation system used to translate the sentences. **Competition** refers to the quality estimation competition in which the data was released and the last column indicates the number of instances the train dataset has for each language pair respectively.

4 Experimental Setup

4.1 QE Dataset

We used several language pairs for which word-level QE annotations were available: English-Chinese (En-Zh), English-Czech (En-Cs), English-German (En-De), English-Russian (En-Ru), English-Latvian (En-Lv) and German-English (De-En). The texts are from a variety of domains and the translations were produced using both neural and statistical machine translation systems. More details about these datasets can be found in Table 1 and in (Specia et al., 2018; Fonseca et al., 2019; Specia et al., 2020).

4.2 Evaluation Criteria

For evaluation, we used the approach proposed in the WMT shared tasks in which the classification performance is calculated using the multiplication of F1-scores for the ‘OK’ and ‘BAD’ classes against the true labels independently: words in the target (‘OK’ for correct words, ‘BAD’ for incorrect words), gaps in the target (‘OK’ for genuine gaps, ‘BAD’ for gaps indicating missing words) and source words (‘BAD’ for words that lead to errors in the target, ‘OK’ for other words) (Specia et al., 2018). In recent WMT shared tasks, the most popular category was predicting quality for words in the target. Therefore, in Section 5 we only report the F1-score for words in the target. Other results are presented in the supplementary material. Prior to WMT 2019, organisers provided separate scores for gaps and words in the target, while after WMT 2019 they produce a single result for target gaps

and words. We follow this latter approach.

5 Results

The values displayed diagonally across section I of Table 2 show the results for supervised, bilingual, word-level QE models where the model was trained on the training set of a particular language pair and tested on the test set of the same language pair. As can be seen in section V, the architecture outperforms the baselines in all the language pairs and also outperforms the majority of the best systems from previous competitions. In addition to the target word F1-score, our architecture outperforms the baselines and best systems in target gaps F1-score and source words F1-score too as shown in Tables 5 and 6. In the following sections we explore its behaviour in different multilingual settings.

5.1 Multilingual QE

We combined instances from all the language pairs and built a single word-level QE model. Our results, displayed in section II (“All”) of Table 2, show that multilingual models perform on par with bilingual models or even better for some language pairs. We also investigate whether combining language pairs that share either the same domain or MT type can be more beneficial, since it is possible that the learning process is better when language pairs share certain characteristics. However as shown in sections III and IV of Table 2, for the majority of the language pairs, specialised multilingual models built on certain domains or MT types do not perform better than multilingual models which contain all the data.

Train Language(s)	IT				Pharmaceutical			Wiki	
	En-Cs SMT	En-De NMT	En-De SMT	En-Ru NMT	De-En SMT	En-LV NMT	En-Lv SMT	En-De NMT	En-Zh NMT
En-Cs SMT	0.6081	(-0.09)	(-0.07)	(-0.09)	(-0.15)	(-0.02)	(-0.01)	(-0.10)	(-0.11)
En-De NMT	(-0.17)	0.4421	(-0.06)	(-0.02)	(-0.18)	(-0.01)	(-0.02)	(-0.01)	(-0.08)
En-De SMT	(-0.01)	(-0.05)	0.6348	(-0.67)	(-0.14)	(-0.06)	(-0.04)	(-0.06)	(-0.09)
En-Ru NMT	(-0.14)	(-0.08)	(-0.16)	0.5592	(-0.12)	(-0.01)	(-0.03)	(-0.09)	(-0.08)
I De-En SMT	(-0.43)	(-0.23)	(-0.33)	(-0.31)	0.6485	(-0.29)	(-0.32)	(-0.25)	(-0.28)
En-LV NMT	(-0.12)	(-0.09)	(-0.14)	(-0.03)	(-0.12)	0.5868	(-0.01)	(0.09)	(-0.08)
En-Lv SMT	(-0.04)	(-0.16)	(-0.10)	(-0.09)	(-0.16)	(-0.01)	0.5939	(-0.15)	(-0.14)
En-De NMT	(-0.11)	(-0.01)	(-0.08)	(-0.02)	(-0.14)	(-0.02)	(-0.04)	0.5013	(-0.06)
En-Zh NMT	(-0.19)	(-0.08)	(-0.17)	(-0.03)	(-0.16)	(-0.03)	(-0.06)	(-0.07)	0.5402
II All	0.6112	0.4523	0.6583	0.5558	0.6221	0.5991	0.5980	0.5101	0.5229
All-1	(-0.01)	(-0.01)	(-0.05)	(-0.02)	(-0.12)	(-0.01)	(-0.01)	(-0.01)	(-0.05)
III Domain	0.6095	0.4467	0.6421	0.5560	0.6331	0.5892	0.5951	0.5021	0.5210
IV SMT/NMT	0.6092	0.4461	0.6410	0.5421	0.6320	0.5885	0.5934	0.5010	0.5205
V Baseline-Marmot	0.4449	0.1812	0.3630	NR	0.4373	0.4208	0.3445	NR	NR
Baseline-OpenKiwi	NR	NR	NR	0.2412	NR	NR	NR	0.4111	0.5583
Best system	0.4449	0.4361	0.6246	0.4780	0.6012	0.4293	0.3618	0.6186	0.6415

Table 2: Target F1-Multi between the algorithm predictions and human annotations. Best results for each language by any method are marked in bold. Sections I, II and III indicate the different evaluation settings. Section IV shows the results of the state-of-the-art methods and the best system submitted for the language pair in that competition. **NR** implies that a particular result was *not reported* by the organisers. Zero-shot results are coloured in grey and the value shows the difference between the best result in that section for that language pair and itself.

5.2 Zero-shot QE

To test whether a QE model trained on a particular language pair can be generalised to other language pairs, different domains and MT types, we performed zero-shot quality estimation. We used the QE model trained on a particular language pair and evaluated it on the test sets of the other language pairs. Non-diagonal values of section I in Table 2 show how each QE model performed on other language pairs. For better visualisation, the non-diagonal values of section I of Table 2 show by how much the score changes when the zero-shot QE model is used instead of the bilingual QE model. As can be seen, the scores decrease, but this decrease is negligible and is to be expected. For most pairs, the QE model that did not see any training instances of that particular language pair outperforms the baselines that were trained extensively on that particular language pair. Further analysing the results, we can see that zero-shot QE performs better when the language pair shares some properties such as domain, MT type or language direction. For example, En-De SMT \Rightarrow En-Cs SMT is better than En-De NMT \Rightarrow En-Cs SMT and En-De SMT \Rightarrow En-De NMT is better than En-Cs SMT \Rightarrow En-De NMT.

We also experimented with zero-shot QE with multilingual QE models. We trained the QE model in all the pairs except one and performed predic-

tion on the test set of the language pair left out. In section II (“All-1”), we show its difference to the multilingual QE model. This also provides competitive results for the majority of the languages, proving it is possible to train a single multilingual QE model and extend it to a multitude of languages and domains. This approach provides better results than performing transfer learning from a bilingual model.

One limitation of the zero-shot QE is its inability to perform when the language direction changes. In the scenario where we performed zero-shot learning from De-En to other language pairs, results degraded considerably from the bilingual result. Similarly, the performance is rather poor when we test on De-En for the multilingual zero-shot experiment as the direction of all the other pairs used for training is different. This is in line with results reported by [Ranasinghe et al. \(2020b\)](#) for sentence level.

5.3 Few-shot QE

We also evaluated how the QE models behave with a limited number of training instances. For each language pair, we initiated the weights of the bilingual model with those of the relevant All-1 QE and trained it on 100, 200, 300 and up to 1000 training instances. We compared the results with those obtained having trained the QE model from scratch for that language pair. The results in Figure 2 show

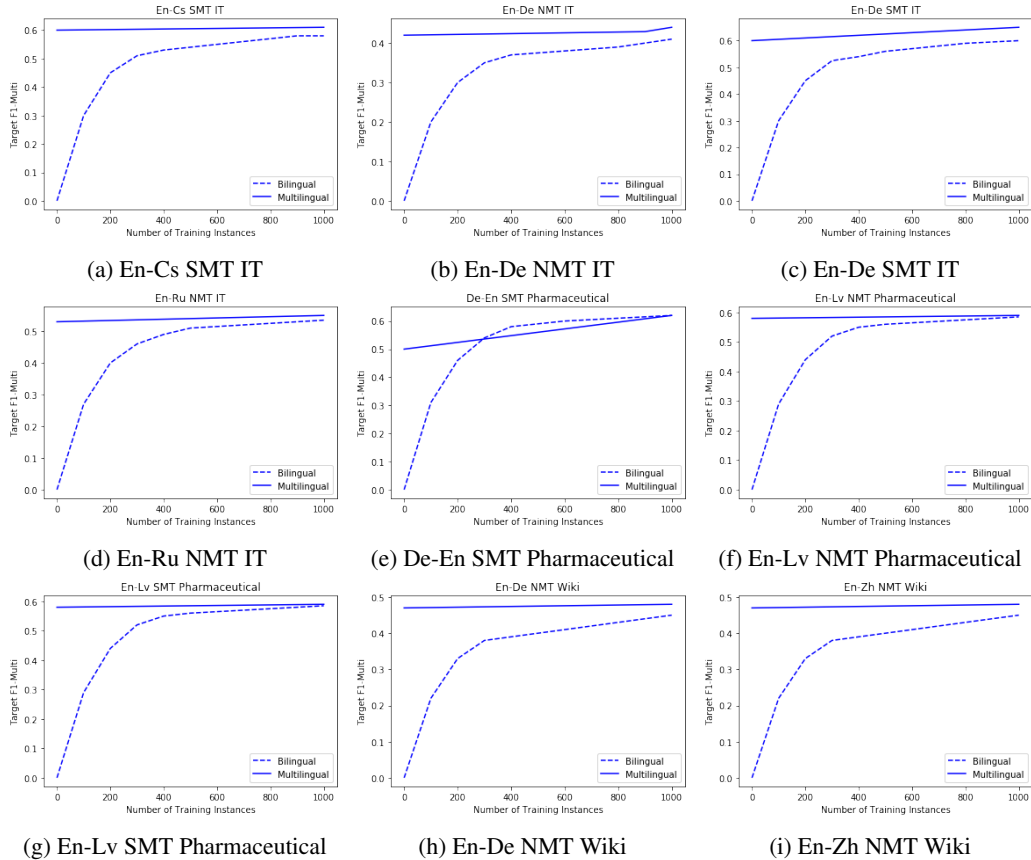


Figure 2: Target F1-Multi scores with Few-shot learning

that All-1 or the multilingual model performs well above the QE model trained from scratch (Bilingual) when there is a limited number of training instances available. Even for the De-En language pair, for which we had comparatively poor zero-shot results, the multilingual model provided better results with a few training instances. It seems that having the model weights already fine-tuned in the multilingual model provides an additional boost to the training process which is advantageous in a few-shot scenario.

6 Conclusions

In this paper, we explored multilingual, word-level QE with transformers. We introduced a new architecture based on transformers to perform word-level QE. The implementation of the architecture, which is based on Hugging Face (Wolf et al., 2020), has been integrated into the TransQuest framework (Ranasinghe et al., 2020b) which won the WMT 2020 QE task (Specia et al., 2020) on sentence-level direct assessment (Ranasinghe et al., 2020a)².

²TransQuest is available on GitHub <https://github.com/tharindudr/TransQuest>

In our experiments, we observed that multilingual QE models deliver excellent results on the language pairs they were trained on. In addition, the multilingual QE models perform well in the majority of the zero-shot scenarios where the multilingual QE model is tested on an unseen language pair. Furthermore, multilingual models perform very well with few-shot learning on an unseen language pair when compared to training from scratch for that language pair, proving that multilingual QE models are effective even with a limited number of training instances. While we centered our analysis around the F1-score of the target words, these findings are consistent with the F1-score of the target gaps and the F1-score of the source words too. This suggests that we can train a single multilingual QE model on as many languages as possible and apply it on other language pairs as well. These findings can be beneficial to perform QE in low-resource languages for which the training data is scarce and when maintaining several QE models for different language pairs is arduous.

References

- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Vishrav Chaudhary, Mark Fishel, Francisco Guzmán, and Lucia Specia. 2020. [BERGAMOT-LATTE submissions for the WMT20 quality estimation shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1010–1017, Online. Association for Computational Linguistics.
- Erick Fonseca, Lisa Yankovskaya, André F. T. Martins, Mark Fishel, and Christian Federmann. 2019. [Findings of the WMT 2019 shared tasks on quality estimation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–10, Florence, Italy. Association for Computational Linguistics.
- Julia Ive, Frédéric Blain, and Lucia Specia. 2018. [deep-Quest: A framework for neural-based quality estimation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3146–3157, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, and André F. T. Martins. 2019. [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.
- Dongjun Lee. 2020. [Two-phase cross-lingual language model fine-tuning for machine translation quality estimation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1024–1028, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Varvara Logacheva, Chris Hokamp, and Lucia Specia. 2016. [MARMOT: A toolkit for translation quality estimation at the word level](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3671–3674, Portorož, Slovenia. European Language Resources Association (ELRA).
- Toan Q. Nguyen and David Chiang. 2017. [Transfer learning across low-resource, related languages for neural machine translation](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 296–301, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- 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 and Marcos Zampieri. 2020. [Multilingual offensive language identification with cross-lingual embeddings](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5838–5844, Online. Association for Computational Linguistics.
- Tharindu Ranasinghe and Marcos Zampieri. 2021. [MUDes: Multilingual detection of offensive spans](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 144–152, Online. Association for Computational Linguistics.
- Kashif Shah and Lucia Specia. 2016. [Large-scale multitask learning for machine translation quality estimation](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 558–567, San Diego, California. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán,

- and André F. T. Martins. 2020. [Findings of the WMT 2020 shared task on quality estimation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- 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.
- Lucia Specia, Marco Turchi, Nicola Cancedda, Nello Cristianini, and Marc Dymetman. 2009. [Estimating the sentence-level quality of machine translation systems](#). In *Proceedings of the 13th Annual conference of the European Association for Machine Translation*, Barcelona, Spain. European Association for Machine Translation.
- Shuo Sun, Marina Fomicheva, Frédéric Blain, Vishrav Chaudhary, Ahmed El-Kishky, Adithya Renduchintala, Francisco Guzmán, and Lucia Specia. 2020. [An exploratory study on multilingual quality estimation](#). 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 366–377, Suzhou, China. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Trans-formers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.