

TransQuest: Translation 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

Recent years have seen big advances in the field of sentence-level quality estimation (QE), largely as a result of using neural-based architectures. However, the majority of these methods work only on the language pair they are trained on and need retraining for new language pairs. This process can prove difficult from a technical point of view and is usually computationally expensive. In this paper we propose a simple QE framework based on cross-lingual transformers, and we use it to implement and evaluate two different neural architectures. Our evaluation shows that the proposed methods achieve state-of-the-art results outperforming current open-source quality estimation frameworks when trained on datasets from WMT. In addition, the framework proves very useful in transfer learning settings, especially when dealing with low-resourced languages, allowing us to obtain very competitive results.

1 Introduction

The goal of quality estimation (QE) is to evaluate the quality of a translation without having access to a reference translation (Specia et al., 2018). High-accuracy QE that can be easily deployed for a number of language pairs is the missing piece in many commercial translation workflows as they have numerous potential uses. They can be employed to select the best translation when several translation engines are available or can inform the end user about the reliability of automatically translated content. In addition, QE systems can be used to decide whether a translation can be published as it is in a given context, or whether it requires human post-editing before publishing or even translation from scratch by a human (Kepler et al., 2019). The estimation of translation quality can be done at different levels: document level, sentence level and word/phrase level (Ive et al., 2018). In this research we focus on sentence-level quality estimation.

As we discuss in Section 2, at present neural-based QE methods constitute the state of the art in quality estimation. However, these approaches are based on complex neural networks and require resource-intensive training. This resource-intensive nature of these deep-learning-based frameworks makes it expensive to have QE systems that work for several languages at the same time. Furthermore, these architectures require a large number of annotated instances for training, making the quality estimation task very difficult for low-resource language pairs.

In this paper we propose *TransQuest*, a framework for sentence-level machine translation quality estimation which solves the aforementioned problems, whilst obtaining competitive results. The motivation behind this research is to propose a simple architecture which can be easily trained with different types of inputs (i.e. different language pairs or language from different domains) and can be used for transfer learning in settings where there is not enough training data. We show that *TransQuest* outperforms current open-source quality estimation frameworks and compares favourably to winning solutions submitted to recent shared tasks on 15 different language pairs on different aspects of quality estimation. In fact, a tuned version of TransQuest was declared the winner for all 8 tasks of the direct

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

assessment sentence level QE shared task organised at WMT 2020 (for more details see (Ranasinghe et al., 2020) and Section 5.1. The main contributions of this paper are the following:

1. We introduce *TransQuest*, an open-source framework, and use it to implement two neural network architectures that outperform current state-of-the-art quality estimation methods in two different aspects of sentence-level quality estimation.
2. To the best of our knowledge this is the first neural-based method which develops a model capable of providing quality estimation for more than one language pair. In this way we address the problem of high costs required to maintain a multi-language-pair QE environment.
3. We tackle the problem of quality estimation in low-resource language pairs by showing that even with a small number of annotated training instances, *TransQuest* with transfer learning can outperform current state-of-the-art quality estimation methods in low-resource language pairs.
4. We provide important resources to the community: the code as an open-source framework, as well as the *TransQuest* model zoo – a collection of pre-trained quality estimation models which have been trained on 15 different language pairs and different aspects of quality estimation – will be freely available to the community¹.

The remainder of the paper is structured as follows. We first present a brief overview of related work in order to define the context of our work. In Section 3 we present the TransQuest framework and the methodology employed to train it. The datasets used to train it are presented in Section 4, followed by the evaluation and discussion in Section 5. The paper finishes with conclusions and ideas for future research directions.

2 Related Work

During the past decade there has been tremendous progress in the field of quality estimation, largely as a result of the QE shared tasks organised annually by the Workshops on Statistical Machine Translation (WMT), more recently called the Conferences on Machine Translation, since 2012. The annotated datasets these shared tasks released each year have led to the development of many open-source QE systems like QuEst (Specia et al., 2013), QuEst++ (Specia et al., 2015), deepQuest (Ive et al., 2018), and OpenKiwi (Kepler et al., 2019). Before the neural network era, most of the quality estimation systems like QuEst (Specia et al., 2013) and QuEst++ (Specia et al., 2015) were heavily dependent on linguistic processing and feature engineering to train traditional machine-learning algorithms like support vector regression and randomised decision trees (Specia et al., 2013). Even though, they provided good results, these traditional approaches are no longer the state of the art. In recent years, neural-based QE systems have consistently topped the leader boards in WMT quality estimation shared tasks (Kepler et al., 2019).

For example, the best-performing system at the WMT 2017 shared task on QE was POSTECH, which is purely neural and does not rely on feature engineering at all (Kim et al., 2017). POSTECH revolves around an encoder-decoder Recurrent Neural Network (RNN) (referred to as the 'predictor'), stacked with a bidirectional RNN (the 'estimator') that produces quality estimates. In the predictor, an encoder-decoder RNN model predicts words based on their context representations and in the estimator step there is a bidirectional RNN model to produce quality estimates for words, phrases and sentences based on representations from the predictor. To be effective, POSTECH requires extensive predictor pre-training, which means it depends on large parallel data and is computationally intensive (Ive et al., 2018). The POSTECH architecture was later re-implemented in deepQuest (Ive et al., 2018).

OpenKiwi (Kepler et al., 2019) is another open-source QE framework developed by Unbabel. It implements four different neural network architectures QUETCH (Kreutzer et al., 2015), NUQE (Martins et al., 2016), Predictor-Estimator (Kim et al., 2017) and a stacked model of those architectures. Both the QUETCH and NUQE architectures have simple neural network models that do not rely on

¹The public GitHub repository is available on <https://github.com/tharindudr/transquest> and the official documentation is available on <https://tharindudr.github.io/TransQuest>.

additional parallel data, but do not perform that well. The Predictor-Estimator model is similar to the POSTECH architecture and relies on additional parallel data. In OpenKiwi, the best performance for sentence-level quality estimation was given by the stacked model that used the Predictor-Estimator model, meaning that the best model requires extensive predictor pre-training and relies on large parallel data and computational resources.

In order to remove the dependency on large parallel data, which also entails the need for powerful computational resources, we propose to use crosslingual embeddings that are already fine-tuned to reflect properties between languages. We assume that by using them we will ease the burden of having complex neural network architectures. Over the last few years there has been significant work done in the area of crosslingual embeddings (Ruder et al., 2019).

Since the introduction of BERT (Devlin et al., 2019), transformer models have been used successfully for various NLP tasks such as named entity recognition (Devlin et al., 2019), sentence classification (Sun et al., 2019), and question answering (Devlin et al., 2019), in many cases improving the state of the art. Most of the tasks were focused on English due to the fact that most of the pre-trained transformer models were trained on English data. Although there are several multilingual models like multilingual BERT (mBERT) (Devlin et al., 2019) and multilingual DistilBERT (mDistilBERT) (Sanh et al., 2019), researchers expressed some reservations about their ability to represent all the languages (Pires et al., 2019). In addition, although mBERT and mDistilBERT showed some crosslingual characteristics, they do not perform well on crosslingual benchmarks (K et al., 2020).

XLM-RoBERTa (XLM-R) was released in November 2019 (Conneau et al., 2020) as an update to the XLM-100 model (Conneau and Lample, 2019). XLM-R takes a step back from XLM, eschewing XLM’s Translation Language Modeling (TLM) objective since it requires a dataset of parallel sentences, which can be difficult to acquire. Instead, XLM-R trains RoBERTa (Liu et al., 2019) on a huge, multilingual dataset at an enormous scale: unlabelled text in 104 languages is extracted from CommonCrawl datasets, totalling 2.5TB of text. It is trained using only RoBERTa’s (Liu et al., 2019) masked language modelling (MLM) objective. Surprisingly, this strategy provided better results in crosslingual tasks. XLM-R outperforms mBERT on a variety of crosslingual benchmarks such as crosslingual natural language inference and crosslingual question answering (Conneau et al., 2020).

Both architectures proposed in *TransQuest* have been successfully applied in the monolingual semantic textual similarity tasks (Devlin et al., 2019; Reimers and Gurevych, 2019). When applied in monolingual experiments, both of them use monolingual transformer models like BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) as the input. This inspired us to change the input in such a way that it can represent both the source and target sentences for which the quality of translation needs to be estimated, with the hope that the same architectures would also provide good results in the QE task. Our initial experiments showed that crosslingual embeddings like XLM-R provide better results than multilingual embeddings like mBERT. Therefore, in this research we explore the performance of crosslingual embeddings with simple neural network architectures for the sentence-level quality estimation task. To the best of our knowledge, state-of-the-art crosslingual contextual embeddings such as XLM-R have not been used in quality estimation before.

3 Methodology

This section presents the methodology used to develop our quality estimation methods. We first describe the neural network architectures we proposed, followed by the method used to train these architectures.

3.1 Neural Network Architectures

The *TransQuest* framework that is used to implement the two architectures described here relies on the XLM-R transformer model (Conneau et al., 2020) to derive the representations of the input sentences. The XLM-R transformer model takes a sequence of no more than 512 tokens as input and outputs the representation of the sequence. The first token of the sequence is always [CLS], which contains the special embedding to represent the whole sequence, followed by embeddings acquired for each word in the sequence. As shown below, our proposed neural network architectures can utilise both the embedding

for the [CLS] token and the embeddings generated for each word. The output of the transformer (or transformers for **SiameseTransQuest** described below), is fed into a simple output layer which is used to estimate the quality of translation. We describe below the way the XLM-R transformer is used and the output layer, as they are different in the two instantiations of the framework. The fact that we do not rely on a complex output layer makes training our architectures much less computational intensive than alternative solutions. The *TransQuest* framework is open-source, which means researchers can easily propose alternative architectures to the ones we present in this paper.

Both neural network architectures presented below use the pre-trained XLM-R-large model released by HuggingFace’s Transformers library (Wolf et al., 2019). The XLM-R-large model covers 104 languages (Conneau et al., 2020), making it potentially very useful to estimate the translation quality for a large number of language pairs.

TransQuest implements two different neural network architectures to perform sentence-level translation quality estimation which we describe below. The architectures are presented in Figure 1.

1. **MonoTransQuest (MTransQuest)**: The first architecture proposed uses a single XLM-R transformer model and is shown in Figure 1a. The input of this model is a concatenation of the original sentence and its translation, separated by the [SEP] token. We experimented with three pooling strategies for the output of the transformer model: using the output of the [CLS] token (CLS-strategy); computing the mean of all output vectors of the input words (MEAN-strategy); and computing a max-over-time of the output vectors of the input words (MAX-strategy). The output of the pooling strategy is used as the input of a softmax layer that predicts the quality score of the translation. We used mean-squared-error loss as the objective function. Early experiments we carried out demonstrated that the CLS-strategy leads to better results than the other two strategies for this architecture. Therefore, we used the embedding of the [CLS] token as the input of a softmax layer.
2. **SiameseTransQuest (STransQuest)**: The second approach proposed in this paper relies on the Siamese architecture depicted in Figure 1b which has shown promising results in monolingual semantic textual similarity tasks (Reimers and Gurevych, 2019; Ranasinghe et al., 2019). In this case, we feed the original text and the translation into two separate XLM-R transformer models. Similar to the previous architecture we used the same three pooling strategies for the outputs of the transformer models. We then calculated the cosine similarity between the two outputs of the pooling strategy. We used mean-squared-error loss as the objective function. In initial experiments we carried out with this architecture, the MEAN-strategy showed better results than the other two strategies. For this reason, we used the MEAN-strategy for our experiments. Therefore, cosine similarity is calculated between the the mean of all output vectors of the input words produced by each transformer.

3.2 Training Details

We used the same set of configurations for all the language pairs evaluated in this paper in order to ensure consistency between all the languages. This also provides a good starting configuration for researchers who intend to use *TransQuest* on a new language pair. In both architectures we used a batch-size of eight, Adam optimiser with learning rate $2e-5$, and a linear learning rate warm-up over 10% of the training data. During the training process, the parameters of XLM-R model, as well as the parameters of the subsequent layers, were updated. The models were trained using only training data. Furthermore, they were evaluated while training using an evaluation set that had one fifth of the rows in training data. We performed early stopping if the evaluation loss did not improve over ten evaluation steps. All the models were trained for three epochs. For some of the experiments, we used an Nvidia Tesla K80 GPU, whilst for others we used an Nvidia Tesla T4 GPU. This was purely based on the availability of the hardware and it was not a methodological decision.

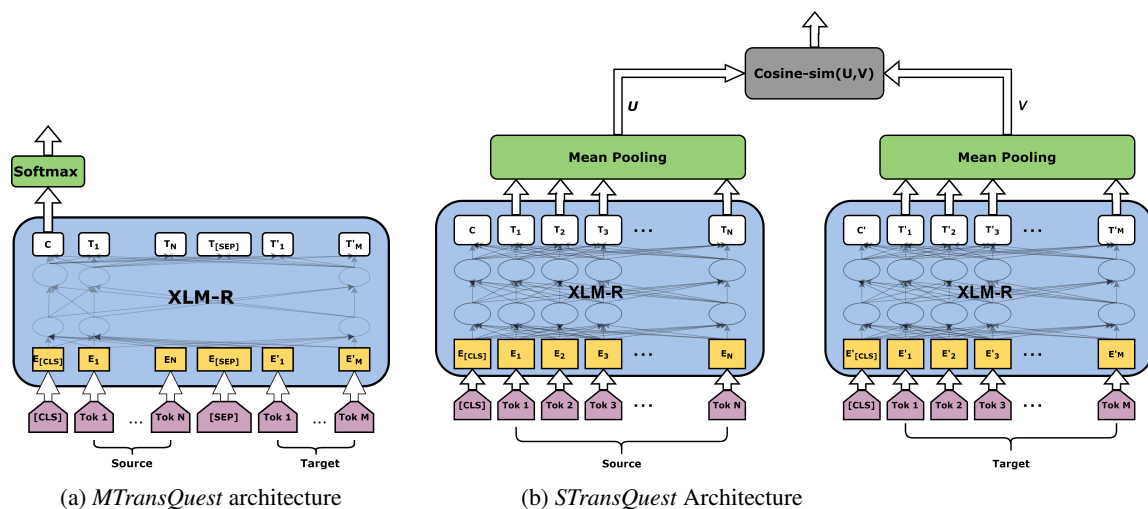


Figure 1: Two architectures of the *TransQuest* framework.

4 Dataset

We used the architectures described above to predict two standard measures that express the quality of a translation: Human-mediated Translation Edit Rate (HTER) and Direct Assessment (DA). All the datasets that we used are publicly available and were released in WMT quality estimation tasks in recent years (Specia et al., 2018; Fonseca et al., 2019; Specia et al., 2020). This was done to ensure replicability of our experiments and to allow us to compare our results with the state of the art. In the remainder of the section we provide more details about the datasets used.

4.1 Predicting HTER

The performance of QE systems has typically been assessed using the semiautomatic HTER (Human-mediated Translation Edit Rate). HTER is an edit-distance-based measure which captures the distance between the automatic translation and a reference translation in terms of the number of modifications required to transform one into another. In light of this, a QE system should be able to predict the percentage of edits required in the translation. We used several language pairs for which HTER information was 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).

4.2 Predicting DA

Even though HTER has been typically used to assess quality in machine translations, the reliability of this metric for assessing the performance of quality estimation systems has been questioned by researchers (Graham et al., 2016). The current practice in MT evaluation is the so-called Direct Assessment (DA) of MT quality (Graham et al., 2017), where raters evaluate the machine translation on a continuous 1-100 scale. This method has been shown to improve the reproducibility of manual evaluation and to provide a more reliable gold standard for automatic evaluation metrics (Graham et al., 2015).

We used a recently created dataset to predict DA in machine translations which was released for the WMT 2020 quality estimation shared task 1 (Specia et al., 2020). The dataset is composed of data extracted from Wikipedia for six language pairs, consisting of high-resource English-German (En-De) and English-Chinese (En-Zh), medium-resource Romanian-English (Ro-En) and Estonian-English (Et-En), and low-resource Sinhala-English (Si-En) and Nepalese-English (Ne-En), as well as a Russian-

²Language codes are available in ISO 639-1 Registration Authority Website Online - https://www.loc.gov/standards/iso639-2/php/code_list.php

Language Pair	Source	MT system	Competition	train, dev, test size
De-En	Pharmaceutical	Phrase-based SMT	WMT 2018	25,963, 1,000, 1,000
En-Zh	Wiki	fairseq based NMT	WMT 2020	7,000, 1,000, 1,000
En-Cs	IT	Phrase-based SMT	WMT 2018	40,254, 1,000, 1,000
En-De	IT	fairseq based NMT	WMT 2019	13,442, 1,000, 1,000
En-De	IT	Phrase-based SMT	WMT 2018	26,273, 1,000, 1,000
En-Ru	Tech	Online NMT	WMT 2019	15,089, 1,000, 1,000
En-Lv	Pharmaceutical	Attention-based NMT	WMT 2018	12,936, 1,000, 1,000
En-Lv	Pharmaceutical	Phrase-based SMT	WMT 2018	11,251, 1,000, 1,000

Table 1: Information about language pairs used to predict HTER. The **Language Pair** column shows the language pairs we used in ISO 639-1 codes². **Source** expresses the domain of the sentence and **MT system** is the Machine Translation system used to translate the sentences. In that column NMT indicates Neural Machine Translation and SMT indicates Statistical Machine Translation. **Competition** shows the quality estimation competition in which the data was released and the last column indicates the number of instances the train, development and test dataset had in each language pair respectively.

English (En-Ru) dataset which combines articles from Wikipedia and Reddit (Fomicheva et al., 2020). These datasets have been collected by translating sentences sampled from source-language articles using state-of-the-art NMT models built using the fairseq toolkit (Ott et al., 2019) and annotated with DA scores by professional translators. Each translation was rated with a score from 0-100 according to the perceived translation quality by at least three translators (Specia et al., 2020). The DA scores were standardised using the z-score. The quality estimation systems evaluated on these datasets have to predict the mean DA z-scores of test sentence pairs. Each language pair has 7,000 sentence pairs in the training set, 1,000 sentence pairs in the development set and another 1,000 sentence pairs in the testing set.

5 Evaluation and discussion

This section presents the evaluation results of our architectures on the datasets described in the previous section in a variety of settings. We first evaluate them in a single language pair setting (Section 5.1), which is essentially the setting employed in the WMT shared tasks. We then evaluate in a setting where we combine datasets in several languages for training (Section 5.2). We conclude the section with an evaluation of a transfer learning setting (Section 5.3).

In order to better understand the performance of our approach, we compare our results with the baselines reported by the WMT2018-2020 organisers. Two baselines were used: OpenKiwi (Kepler et al., 2019) and QuEst++ (Specia et al., 2015). Row **IV** of Tables 2 and 3 present the results for these baselines. For some language pairs in the 2018 WMT quality estimation shared task the organisers did not report the scores obtained using OpenKiwi. These cases are marked with **NR** in Table 2. Row **IV** of both tables also includes the results of the best system from WMT2018 to WMT2020 for each setting. Additionally, Row **IV** of Table 3 shows the results of the TransQuest’s submission to WMT 2020 QE Task 1, which was the winning solution in all the languages (Ranasinghe et al., 2020).

The evaluation metric used was the Pearson correlation (r) between the predictions and the gold standard from the test set, which is the most commonly used evaluation metric in recent WMT quality estimation shared tasks (Specia et al., 2018; Fonseca et al., 2019; Specia et al., 2020).

5.1 Supervised Single Language Pair Quality Estimation

The first evaluation we carried out was the *supervised single language pair* evaluation where we used the training set of each language to build a quality estimation model and we evaluated it on a testing set from the same language. This replicates the standard QE evaluation carried out in the WMT shared tasks. The results for each language in *supervised* settings are shown in row I of Tables 2 and 3. The results indicate that both architectures proposed in *TransQuest* outperform the baselines in all the language pairs of both

Method	Mid-resource				High-resource				
	En-Cs SMT	En-Ru NMT	En-Lv SMT	En-Lv NMT	De-En SMT	En-Zh NMT	En-De SMT	En-De NMT	
I	MTransQuest	0.7207	0.7126	0.6592	0.7394	0.7939	0.6119	0.7137	0.5994
	STransQuest	0.6853	0.6723	0.6320	0.7183	0.7524	0.5821	0.6992	0.5875
II	MTransQuest *-En En-*	0.7168	0.7046	0.7181	0.7482	0.7939	0.6101	0.7355	0.5992
	STransQuest *-En En-*	0.6663	0.6701	0.6533	0.7192	0.7524	0.5721	0.7000	0.5793
III	MTransQuest-m	0.7111	0.7012	0.7141	0.7450	0.7878	0.6092	0.7300	0.5982
	STransQuest-m	0.6561	0.6614	0.6621	0.7202	0.7369	0.5612	0.7015	0.5771
IV	Quest ++	0.3943	0.2601	0.3528	0.4435	0.3323	NR	0.3653	NR
	OpenKiwi	NR	0.5923	NR	NR	NR	0.5058	0.7108	0.4001
	Best system	0.6918	0.5923	0.6188	0.6819	0.7888	0.6641	0.7397	0.5718
V	mBERT	0.6423	0.6354	0.5772	0.6531	0.7005	0.5483	0.6239	0.5002

Table 2: Pearson (r) correlation between *TransQuest* algorithm predictions and human post-editing effort. Best results for each language by any method are marked in bold. Rows I, II and III indicate the different evaluation settings, explained in Sections 5.1, 5.2 and 5.3. Row 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. Row V presents the results of the multilingual BERT (mBERT) model in MonoTransQuest Architecture.

aspects in quality estimation, and also outperform the best systems from previous competitions. From the two architectures, *MTransQuest* performs slightly better than *STransQuest*.

In the HTER aspect of quality estimation, as shown in Table 2, *MTransQuest* gains ≈ 0.1 - 0.2 Pearson correlation boost over OpenKiwi in most language pairs. However, OpenKiwi comes very close to *MTransQuest* in En-De SMT. In the language pairs where OpenKiwi results are not available *MTransQuest* gains ≈ 0.3 - 0.4 Pearson correlation boost over QuEst++ in all language pairs for both NMT and SMT. Table 2 also gives the results of the best system submitted for a particular language pair. It is worth noting that for the training setting described in this section, the *TransQuest* results surpass the best system in all the language pairs with the exception of the En-De SMT and En-Zh NMT datasets.

As shown in Table 3, in the DA aspect of quality estimation, *MTransQuest* gained ≈ 0.2 - 0.3 Pearson correlation boost over OpenKiwi in all the language pairs. Additionally, *MTransQuest* achieves ≈ 0.4 Pearson correlation boost over OpenKiwi in the low-resource language pair Ne-En. Furthermore, TransQuest participated in WMT 2020 quality estimation shared task 1 and it was the winning solution in all the language pairs. To achieve this result, TransQuest was fine-tuned with self-ensemble and data augmentation to achieve the first place. We do not describe here the fine tuning approaches since they are task specific, but more details can be found in (Ranasinghe et al., 2020).

Additionally, row V in both Tables 2 and 3 shows the results of multilingual BERT (mBERT) in MonoTransQuest architecture. We used the same settings similar to XLM-R. The results show that XLM-R model outperforms the mBERT model in all the language pairs of both aspects in quality estimation and we can safely assume that the cross lingual nature of the XLM-R transformers had a clear impact to the results.

5.2 Supervised Multi-Language Pair Quality Estimation

Most of the available open-source quality estimation frameworks require maintaining separate machine learning models for each language. This can be very challenging in a practical environment where the systems have to work with 10-20 language pairs. Furthermore, pre-trained neural quality estimation models are large. In a commercial environment, where the quality estimation systems need to do inference on several language pairs, loading all of the pre-trained models from all language pairs would require a lot of Random Access Memory (RAM) space and result in a huge cost.

Method	Low-resource		Mid-resource			High-resource		
	Si-En	Ne-En	Et-En	Ro-En	Ru-En	En-De	En-Zh	
I	MTransQuest	0.6525	0.7914	0.7748	0.8982	0.7734	0.4669	0.4779
	STransQuest	0.5957	0.7081	0.6804	0.8501	0.7126	0.3992	0.4067
II	MTransQuest *-En En-*	0.6528	0.7824	0.7827	0.8868	0.7821	0.4518	0.4334
	STransQuest *-En En-*	0.5968	0.6992	0.6921	0.8432	0.7152	0.3621	0.3812
III	MTransQuest-m	0.6526	0.7581	0.7574	0.8856	0.7521	0.4420	0.4646
	STransQuest-m	0.5970	0.6980	0.6934	0.8426	0.6945	0.3832	0.3900
IV	OpenKiwi	0.3737	0.3860	0.4770	0.6845	0.5479	0.1455	0.1902
	TransQuest @WMT2020	0.6849	0.8222	0.8240	0.9082	0.8082	0.5539	0.5373
V	mBERT	NS	0.6452	0.6231	0.8351	0.6661	0.3765	0.3982

Table 3: Pearson (r) correlation between *TransQuest* algorithm predictions and human DA judgments. Best results for each language (any method) are marked in bold. Rows I, II and III indicate the different settings of *TransQuest*, explained in Sections 5.1-5.3. OpenKiwi baseline results are in Row IV. Row IV also shows the results for TransQuest’s submission to WMT 2020 QE Task 1 (Ranasinghe et al., 2020) which was also the winning solution. Row V presents the results of the multilingual BERT (mBERT) model in MonoTransQuest Architecture. **NS** implies that the non-English language in the language pair is not supported by mBERT.

Therefore, with *TransQuest* we propose a single model that can perform quality estimation on several language pairs. We propose two training strategies for *supervised multi language pair* settings:

1. We separate the language pairs into two groups. One group contains all the language pairs where the source language is English and in the other group the target is always English. We represent the former with En-*, and the latter with *-En. We train both architectures by concatenating training sets in all the language pairs in a particular group. To ease the comparison of the results, we evaluate the model separately on each language pair. We do this process for both aspects in quality estimation. The results are shown in row II in Tables 2 and 3.
2. We concatenate training data from all the language pairs, without considering the direction of the translation, and build a single model for all language pairs. We refer to these models by *MTransQuest-m* and *STransQuest-m*. Similarly to the first *multi-language pair* training strategy, we evaluate the model separately on each language pair and for both aspects of quality estimation. The results are shown in row III in Tables 2 and 3.

As depicted in Tables 2 and 3, the multi-language pair experiments yielded very competitive results. In fact, for some language pairs the multi-language pair model performed better than the model that was trained solely on that particular pair. When predicting HTER, the multi-language pair model performed better in En-Lv for both SMT and NMT, and in En-De for SMT while performing on par in En-De for NMT. In predicting DA, the multi-language pair model performed better in Et-En, Ru-En and Si-En. In addition, with the exceptions of the En-De SMT and En-Zh NMT setting, the models that consider the direction of the language pairs are better than the best systems submitted to previous editions of WMT.

Throughout our experiments we noted that the TransQuest models built with the direction of the language pairs in mind performed slightly better than the TransQuest models trained without considering the language pair direction. It should be noted that none of the multi-language pair models’ Pearson correlation decreased by more than 0.03% in any language pair for either TransQuest architecture. Similar to *supervised single language pair* experiments, MTransQuest architecture performed better than STransQuest architecture in all the languages in both aspects of quality estimation.

The size of the pre-trained *TransQuest* models on a single language pair was \approx 2GB. The pre-trained models for multiple language pairs in this section did not exceed more than 2.1 GB. Therefore, we

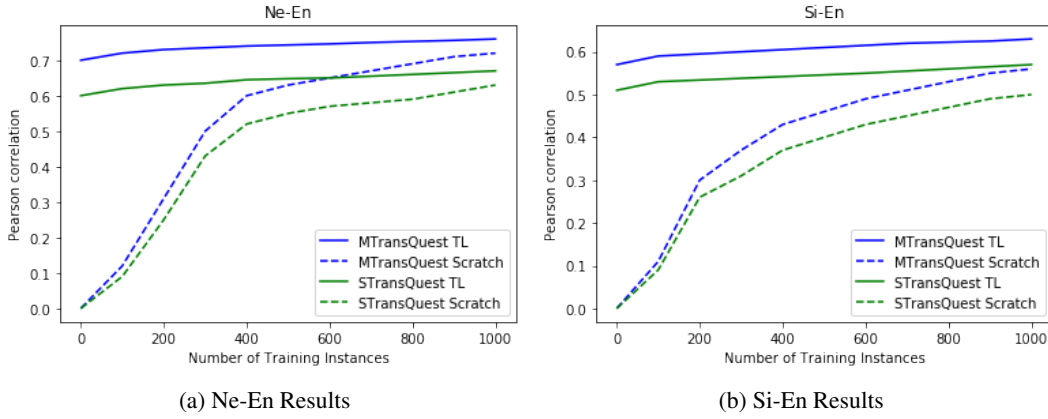


Figure 2: Transfer learning impact of *TransQuest* Architectures. ***MTransQuest Scratch*** and ***STransQuest Scratch*** indicates that the model was trained from scratch while ***MTransQuest TL*** and ***STransQuest TL*** indicates that models followed the transfer learning strategy.

present multi-language pair pre-trained models as a solution for environments that are on tight resources and seek to conduct quality estimation on multiple language pairs.

5.3 Transfer Learning based Quality Estimation

The biggest challenge in building supervised quality estimation models is not having enough annotated data (Fomicheva et al., 2020), especially for low-resourced languages. We explore the possibility of performing transfer learning on low-resource languages using the models trained on better-resourced languages. As the low-resource language pairs were only available in the DA aspect, we conducted this experiment only in the DA aspect in quality estimation. All of the low-resource language pairs in the DA aspect had English as the target language. Considering the positive impact that the direction of the language pair had on Pearson correlation in the previous experiment, we decided to consider only the language pairs with English as the target language. This left us only with mid-resource language pairs for training. We were also aware that the *TransQuest* models had relatively low Pearson correlations with high-resource language pairs in DA.

1. We build a single model for each architecture using all the training data available for mid-resource language pairs: Et-En, Ro-En and Ru-En. We refer to this as *TransQuest-Mid*.
2. When we train a *TransQuest* model for a low-resource language pair, we initiate the model weights from *TransQuest-Mid* and start training. To see whether it is possible to get compatible results even with fewer training instances, we conduct the experiments for 0 (unsupervised), 100, 200, 300 and up to 1,000 training sentence pairs. We do this for Si-En and Ne-En. Depending on the architecture we use, we refer to this model as *MTransQuest TL* or *STransQuest TL*.
3. In order to evaluate the effect of transfer learning we conduct the same experiment in step 2, but we train the model from scratch. Depending on the architecture we use, we refer to this model as *MTransQuest Scratch* or *STransQuest Scratch*.

As shown in Figure 2, the transfer learning strategy significantly impacts the results. In Ne-En, with only 100 training instances, training *MTransQuest scratch* achieves only 0.1242 Pearson correlation between the predictions and gold labels of the test set. However, in Ne-En with only 100 training instances, training *MTransQuest* using the transfer learning strategy achieves 0.7417 Pearson correlation, which is close to the best result obtained with the *MTransQuest* architecture for Ne-En after training with 7,000 instances (0.7914). When the number of training instances grows, the results from the *TransQuest* models trained with the transfer learning strategy and the results from the *TransQuest* models trained from scratch converge. A similar pattern, but with a lower Pearson correlation, can be seen

with *STransQuest*. Similar results can also be observed for the other low-resource language pair, Si-En. Therefore, it is safe to conclude that *TransQuest* with the transfer learning strategy can be hugely beneficial to low-resource language pairs in quality estimation where annotated training instances are scarce.

6 Conclusions

In this paper we introduced *TransQuest*, a new open source framework for quality estimation based on cross-lingual transformers. *TransQuest* is implemented in PyTorch and supports training of sentence-level quality estimation systems on new data. It outperforms other open-source tools on both aspects of sentence-level quality estimation and yields new state-of-the-art quality estimation results. As far as we know, it is the first time that an open-source QE framework has been tested on both aspects of quality estimation. Furthermore, it is the first time that a QE system explores multi-language pair models and transfer learning on low-resource language pairs. Unlike many other open-source neural QE frameworks, *TransQuest* does not use parallel data and hence does not require similar computational resources.

We propose two architectures: *MTransQuest* and *STransQuest*, neither of which have been previously explored in QE tasks. The two architectures have a trade-off between accuracy and efficiency. On an Nvidia Tesla K80 GPU, *MTransQuest* takes 4,480s on average to train on 7,000 instances, while *STransQuest* takes only 3,900s on average for the same experiment. On the same GPU, *MTransQuest* takes 35s on average to perform inference on 1,000 instances which takes *STransQuest* only 16s to do so. Therefore we recommend using *MTransQuest* where accuracy is valued over efficiency, and *STransQuest* where efficiency is prioritised above accuracy. Since there is a growing interest³ in the NLP community for energy efficient machine learning models, we decided to support both architectures in the *TransQuest* Framework.

In the future, we plan to expand *TransQuest* with more neural network architectures and more models for different levels of quality estimation such as word-level and document-level. In the sentence-level, we plan to perform transfer learning on language pairs that do not include English at all. We also hope to conduct unsupervised experiments on low-resource language pairs.

References

- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems 32*, pages 7059–7069. Curran Associates, Inc.
- 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, July. 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, June. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, Vol 8 (2020).
- 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, August. Association for Computational Linguistics.

³Several workshops like *EMC²* (Workshop on Energy Efficient Machine Learning and Cognitive Computing), *SustainNLP 2020* (Workshop on Simple and Efficient Natural Language Processing) has been organised on this aspect.

- Yvette Graham, Timothy Baldwin, and Nitika Mathur. 2015. Accurate evaluation of segment-level machine translation metrics. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1183–1191, Denver, Colorado, May–June. Association for Computational Linguistics.
- Yvette Graham, Timothy Baldwin, Meghan Dowling, Maria Eskevich, Teresa Lynn, and Lamia Tounsi. 2016. Is all that glitters in machine translation quality estimation really gold? In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3124–3134, Osaka, Japan, December. The COLING 2016 Organizing Committee.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2017. Can machine translation systems be evaluated by the crowd alone. *Natural Language Engineering*, 23(1):3–30.
- Julia Ive, Frédéric Blain, and Lucia Specia. 2018. deepQuest: 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, August. Association for Computational Linguistics.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. Cross-lingual ability of multilingual bert: An empirical study. In *International Conference on Learning Representations*.
- 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, July. Association for Computational Linguistics.
- Hyun Kim, Jong-Hyeok Lee, and Seung-Hoon Na. 2017. Predictor-estimator using multilevel task learning with stack propagation for neural quality estimation. In *Proceedings of the Second Conference on Machine Translation*, pages 562–568, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Julia Kreutzer, Shigehiko Schamoni, and Stefan Riezler. 2015. Quality estimation from ScraTCH (QUETCH): Deep learning for word-level translation quality estimation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 316–322, Lisbon, Portugal, September. 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.
- André F. T. Martins, Ramón Astudillo, Chris Hokamp, and Fabio Kepler. 2016. Unbabel’s participation in the WMT16 word-level translation quality estimation shared task. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 806–811, Berlin, Germany, August. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual BERT? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy, July. Association for Computational Linguistics.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2019. Semantic textual similarity with Siamese neural networks. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 1004–1011, Varna, Bulgaria, September. INCOMA Ltd.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020. TransQuest at WMT2020: Sentence-Level Direct Assessment. In *Proceedings of the Fifth Conference on Machine Translation: Shared Task Papers*, online, November. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. 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 3982–3992, Hong Kong, China, November. Association for Computational Linguistics.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. A survey of cross-lingual word embedding models. *J. Artif. Int. Res.*, 65(1):569–630, May.

- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *NeurIPS EM² Workshop*.
- Lucia Specia, Kashif Shah, Jose G.C. de Souza, and Trevor Cohn. 2013. QuEst - a translation quality estimation framework. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 79–84, Sofia, Bulgaria, August. 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, July. Association for Computational Linguistics and The Asian Federation of Natural Language Processing.
- Lucia Specia, Frédéric Blain, Varvara Logacheva, Ramón 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, October. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André FT Martins. 2020. Findings of the wmt 2020 shared task on quality estimation. In *Proceedings of the Fifth Conference on Machine Translation: Shared Task Papers*.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In Maosong Sun, Xuanjing Huang, Heng Ji, Zhiyuan Liu, and Yang Liu, editors, *Chinese Computational Linguistics*, pages 194–206, Cham. Springer International Publishing.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.