

# Verdi: Quality Estimation and Error Detection for Bilingual Corpora

Mingjun Zhao\*  
University of Alberta  
Edmonton, Canada  
zhao2@ualberta.ca

Haijiang Wu\*  
Tencent  
Shenzhen, China  
harywu@tencent.com

Di Niu  
University of Alberta  
Edmonton, Canada  
dniu@ualberta.ca

Zixuan Wang  
Tencent  
Shenzhen, China  
zackiewang@tencent.com

Xiaoli Wang  
Tencent  
Shenzhen, China  
evexlwang@tencent.com

## ABSTRACT

Translation Quality Estimation is critical to reducing post-editing efforts in machine translation and to cross-lingual corpus cleaning. As a research problem, quality estimation (QE) aims to directly estimate the quality of translation in a given pair of source and target sentences, and highlight the words that need corrections, without referencing to golden translations. In this paper, we propose *Verdi*, a novel framework for word-level and sentence-level post-editing effort estimation for bilingual corpora. *Verdi* adopts two word predictors to enable diverse features to be extracted from a pair of sentences for subsequent quality estimation, including a transformer-based neural machine translation (NMT) model and a pre-trained cross-lingual language model (XLM). We exploit the symmetric nature of bilingual corpora and apply model-level dual learning in the NMT predictor, which handles a primal task and a dual task simultaneously with weight sharing, leading to stronger context prediction ability than single-direction NMT models. By taking advantage of the dual learning scheme, we further design a novel feature to directly encode the translated target information without relying on the source context. Extensive experiments conducted on WMT20 QE tasks demonstrate that our method beats the winner of the competition and outperforms other baseline methods by a great margin. We further use the sentence-level scores provided by *Verdi* to clean a parallel corpus and observe benefits on both model performance and training efficiency.

## CCS CONCEPTS

• **Computing methodologies** → **Machine translation**; *Information extraction*; Natural language generation.

## KEYWORDS

Quality Estimation, Machine Translation, Model-level Dual Learning, Bilingual Corpus Filtering

\*Both authors contributed equally to this research.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21, April 19–23, 2021, Ljubljana, Slovenia

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-8312-7/21/04.

<https://doi.org/10.1145/3442381.3449931>

## ACM Reference Format:

Mingjun Zhao, Haijiang Wu, Di Niu, Zixuan Wang, and Xiaoli Wang. 2021. *Verdi: Quality Estimation and Error Detection for Bilingual Corpora*. In *Proceedings of the Web Conference 2021 (WWW '21)*, April 19–23, 2021, Ljubljana, Slovenia. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3442381.3449931>

## 1 INTRODUCTION

Neural Machine Translation (NMT) has achieved great success in generating accurate translations with a single neural network in an end-to-end fashion. Although NMT eliminates the need for hand-crafted features and rules, as compared to traditional statistical machine translation methods, yet a well-performing NMT model is largely dependent on large amounts of high-quality training data. However, acquiring accurate translation pairs is a costly task requiring cross-lingual expertise. Common practices of creating a large translation dataset include leveraging the parallel corpus crawled from web, e.g., Wikipedia, and/or relying on existing machine translators to generate sentence pairs. However, the web-crawled or machine-generated corpora usually do not have any guarantee in quality. They contain a high degree of noise [28] and demand further editing or filtering to correct the errors, which is still a highly expensive procedure. As a result, the ability of automatically estimating the translation quality of sentence pairs becomes an urgent need.

Recently, Quality Estimation (QE) has become of growing importance in the field of machine translation and cross-lingual corpus filtering, which aims to directly assess the quality of a source-target translation pair, without referencing to golden human translations, with typically two goals: word-level and sentence-level post-editing effort estimations [6, 11, 12]. The goal of word-level QE is to predict a tag for each token in the translated sentence to detect the mistranslated words and a tag for each gap between two tokens to detect missing words. On the other hand, in sentence-level QE tasks, a commonly used measurement metric is the Human Translation Error Rate (HTER) score [24], which is the ratio between the number of edits needed and the reference translation length. Table 1 gives a data example from the WMT20 English-Chinese QE dataset, where the accurate translations are shown in green and the bad translations are marked in red.

Traditional studies on quality estimation are based on hand-crafted features [13, 25, 26] or feature selection [1, 22]. Deep learning methods have also been applied to quality estimation [14, 16]. The recent Predictor-Estimator framework [12] for QE has drawn much attention, which consists of an RNN word prediction model firstly trained on a large amount of parallel corpora for feature extractions, and an estimator that is subsequently trained on a smaller amount of sentence pairs with QE labels and provides quality evaluation based on the extracted features. Fan et al. [6] base their work on the predictor-estimator framework and further propose to substitute the RNN predictor with a bidirectional transformer language model.

In this paper, we propose *Verdi*, a novel translation quality estimation and error detection method, which falls into the Predictor-Estimator framework, yet leveraging model-level dual learning, mixture models and cross-lingual language models to extract a range of diverse features to boost performance on both sentence-level and word-level QE tasks. Our framework utilizes two word predictors for feature extraction, including an NMT model and a pre-trained XLM model.

Taking advantage of the symmetric nature of parallel corpora, our NMT word predictor is a transformer [27] reformulated by the model-level dual learning framework [31], which behaves as a regularization method that ties the model parameters that are playing similar roles in two dual tasks: the forward and the backward translation tasks. In particular, we reformulate Transformer into two unified conditional self-attentive encoders that can handle the two dual tasks simultaneously with the same set of parameters, where one conditional encoder works as both the encoder in the source-to-target translation and the decoder in the target-to-source translation. The opposite applies to the other encoder. By training on two dual tasks with double the amount of training data and weight sharing, our NMT predictor is capable of outputting more accurate translations and producing more robust representations for the sentence pair. Additionally, we have designed a novel Dual Model Feature for quality estimation based on the model-level dual learning framework. The dual model feature models the internal correlations in the translated sentence without relying on the source context information, and can effectively improve the performance of quality estimation.

Furthermore, we utilize the mixture models technique in our framework to train the NMT predictor, where multiple expert models are trained and combined. In particular, each expert model is assigned with a different start-of-sentence token, and is trained with the EM algorithm [18], so that different expert models can specialize on different translation styles. Therefore, our NMT predictor is capable of producing diverse translation outputs, thus accommodating the one-to-many setup of machine translation.

The second predictor we utilize is a pre-trained cross-lingual language model (XLM) [4] to further improve the feature extraction ability. Unlike monolingual language models such as BERT [5], XLM models are pre-trained on parallel corpora and construct cross-lingual representations on text, which suit better to the translation quality estimation task. In contrast to the single predictor design in most prior work, e.g., [6, 10–12], our two predictors are capable of modeling the sentence pair from different perspectives, leading

Source	Many butterflies flutter among the flowers and grasses .
MT-back	Many butterflies float in the flowers and grasses .
MT	许多蝴蝶在花草间飘动。
PE (reference)	许多蝴蝶在花草丛中飞舞。
Word tags	OK OK OK OK BAD BAD BAD
Gap Tags	OK OK OK OK OK OK OK OK
HTER Score	0.4286

**Table 1: An example from the WMT20 QE dataset, consisting of the English source sentence, the translated sentence in Chinese (MT) and the back translation in English for demonstration, the reference translation post-edited from MT (PE), and the QE labels including word tags, gap tags and the HTER score.**

to a more compact representation, hence improving the estimation results.

Extensive experiments conducted on the WMT20 QE competition dataset demonstrate that our method beats the winning system of the competition and outperforms the baseline methods by a great margin on both word-level and sentence-level QE tasks. We have also designed and conducted experiments of exploiting the QE estimation results as the filtering criterion for cleaning a bilingual corpus. The outcome shows that compared with using all unfiltered samples, training on the high-quality corpus filtered by our method can yield a better NMT system, as well as save training costs. Our source code is available at <https://github.com/Simpleple/Verdi>.

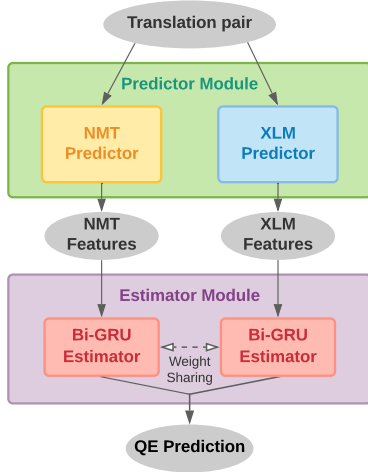
## 2 PROBLEM FORMULATION

In this section, we provide a brief introduction of the background and the problem formulation of the quality estimation for machine translation.

The task of machine translation is to map a sentence  $x$  in the source language to its translation  $y$  in the target language. Modern neural machine translation models adopt an encoder-decoder architecture where an encoder  $p(\mathbf{h}|x)$  maps the source sentence  $x$  into a latent representation  $\mathbf{h}$  and a decoder  $p(y|\mathbf{h}) = \sum_{i=1}^L p(y_i|y_{<i}, \mathbf{h})$  decodes the target sentence  $y$  based on the latent vector  $\mathbf{h}$  and its previous outputs  $\{y_1, \dots, y_{i-1}\}$ , where  $L$  is the length of the target sentence.

Quality estimation aims to provide an evaluation on the quality of the translation pair  $(x, y)$ , either in the detailed word level or in the overall sentence level. In word-level quality estimation, each word in the translated sentence is given a binary tag  $t$  of either “OK” or “BAD” representing the correctness of the translation. Similarly, for each gap in the target sentence, a tag is also provided to indicate whether one or more missing words should be added at the gap position. Therefore, word-level quality estimation can be solved with a typical sequential labeling model  $p(\mathbf{t}|x, y)$  predicting a binary vector  $\mathbf{t}$  for all the words and gaps given the source sentence  $x$  and the translation output  $y$ .

Human Translation Error Rate (HTER) is defined as the ratio between the number of post-editing operations needed to modify



**Figure 1: Illustration of the overall architecture of our proposed framework. It contains a predictor module consisting of an NMT predictor and an XLM predictor for feature extraction, and an estimator module to evaluate the translation quality.**

the translation output to the golden reference and the number of tokens in the translation sentence. The sentence-level quality estimation task aims to predict the HTER score  $l$  for a translation sentence pair and a regression model  $p(l|x, y)$  can be learned to solve it.

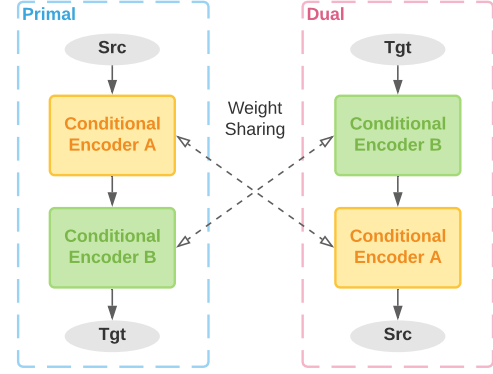
### 3 MODEL DESCRIPTION

In this section, we introduce our proposed framework in detail. Similar to [6, 12], our *Verdi* framework for quality estimation is based on the predictor-estimator architecture, where a predictor module is first trained on a parallel corpus and is responsible for extracting features, and an estimator module is trained on QE datasets to evaluate the translation quality by taking the extracted features as input. Fig.1 illustrates the overall architecture of our framework which mainly consists of three components: the NMT predictor, the XLM predictor, and the QE estimator.

The NMT predictor is a neural machine translation model based on transformer architecture, that automatically transforms sentences in the source language to their translations in the target language. For the purpose of sufficiently exploiting the symmetric structure of machine translation, we propose to adopt model-level dual learning on the NMT predictor, and train it on both directions of translation. Additionally, we apply the method of mixture models [23] for diverse machine translation during the training of the NMT predictor, in order to increase the diversity of the translation outputs and derive more robust features for evaluation.

The XLM predictor utilizes a pre-trained XLM model that takes in both sentences in a translation pair as input, and produces cross-lingual feature representations.

For the design of QE features, we first borrow the idea of quality estimation feature vectors (QEFVs) and mis-matching features from



**Figure 2: Illustration of the training scheme of the NMT predictor with model-level dual learning.**

[6, 12], and propose a novel feature named Dual Model Feature. QEFVs can be seen as the approximated transferred knowledge from word prediction to quality estimation, and mis-matching features measure the difference between the model predictions and the given translated sentence. Apart from the above features, our proposed dual model feature takes the advantage of model-level dual learning, and focuses on modeling the translation information without source contexts involved.

Because the QE datasets are usually small with a few thousand samples, it is very likely to overfit to the dataset. Thus, we incorporate bidirectional GRU estimators with weight sharing to transform the extracted NMT and XLM features into the final quality prediction, to alleviate the overfitting problem.

#### 3.1 NMT Predictor

**3.1.1 Model-level Dual Learning.** Transformer [27] is the current state-of-the-art model in most machine translation tasks, which incorporates the encoder-decoder architecture and is solely built on the multi-headed attention mechanism. Compared to traditional NMT models, such as RNN-based and CNN-based models, the attention mechanism used in transformer can efficiently draw global dependencies between inputs and outputs, and allows significantly more parallelization.

The encoder and decoder of transformer both consist of several stacked blocks. Each block is composed of three types of layers including (1) a multi-headed self-attention layer that takes in a set of hidden states of the previous block and weighs their relevance to each other; (2) an optional multi-headed attention layer over the encodings to draw relevant information from the encoder outputs, which only presents in decoder blocks; and (3) a feed-forward layer that applies non-linear transformations. Between two consecutive layers, a residual connection and a layer normalization operation is performed.

Model-level dual learning [31] can be used to improve the translation model by leveraging the dualities of the machine translation task. It designs a model architecture where a single model can perform both the primal task and the dual task simultaneously with

**Algorithm 1:** Model-level Dual Training of NMT Predictor

---

**Input:** Parallel Corpus  $D$ , an NMT predictor composed of two conditional encoders  $Enc_A$  and  $Enc_B$

```

1 while not converged do
2   sample source and target sentences  $(x, y)$  from  $D$ 
   // the primal task
3   encode  $x$  with  $Enc_A$  to derive  $h_x$ 
4   decode  $h_x$  with  $Enc_B$  to derive predictions  $\tilde{y}$ 
5   update the parameters of  $Enc_A$  and  $Enc_B$  with  $loss(y, \tilde{y})$ 
   // the dual task
6   encode  $x$  with  $Enc_B$  to derive  $h_y$ 
7   decode  $h_y$  with  $Enc_A$  to derive predictions  $\tilde{x}$ 
8   update the parameters of  $Enc_A$  and  $Enc_B$  with  $loss(x, \tilde{x})$ 

```

---

the same set of parameters. In order to apply the model-level dual learning method to the NMT transformer model, we perform some minor modifications to the transformer encoder and decoder, and construct a conditional encoder module that can perform either encoding or decoding.

Specifically, the unified block module has the same model structure with a transformer decoder block, but functions differently in the encoding and decoding stage. When the block works in a decoder, it performs the identical operations with a transformer decoder block. And when the unified block performs encoding functions, instead of the encoder outputs, a tensor of zeros is passed as the  $Q$  and  $V$  input into its second multi-attention layer. The attention to a tensor of zeros will result in an output of zeros as well, which is then restored back to the output of the first self-attention layer through the residual connection. Through these operations, it is guaranteed that the unified conditional encoder will function identically with the traditional transformer encoder and decoder, with a unified structure. The workflow of the  $i$ -th unified block is shown as follows:

$$\begin{aligned}
\mathbf{e}_i &= \text{MultiHead}(K = \mathbf{e}_{i-1}, Q = \mathbf{e}_{i-1}, V = \mathbf{e}_{i-1}), \\
\mathbf{e}_i &= \text{LayerNorm}(\mathbf{e}_i), \\
\mathbf{e}_i &+= \begin{cases} \text{MultiHead}(K = \mathbf{e}_{i-1}, Q = \mathbf{0}, V = \mathbf{0}), & \text{encoding,} \\ \text{MultiHead}(K = \mathbf{e}_{i-1}, Q = \hat{\mathbf{e}}, V = \hat{\mathbf{e}}), & \text{decoding,} \end{cases} \quad (1) \\
\mathbf{e}_i &= \text{LayerNorm}(\mathbf{e}_i), \\
\mathbf{e}_i &+= \text{FeedForward}(\mathbf{e}_i), \\
\mathbf{e}_i &= \text{LayerNorm}(\mathbf{e}_i),
\end{aligned}$$

where  $\mathbf{e}_{i-1}$  is the output of the previous block, and  $\hat{\mathbf{e}}$  is the encoder output required in decoding processes. By simply stacking a few blocks, we can build a unified conditional encoder that can perform exchanging tasks of encoding and decoding.

We build our NMT predictor based on the encoder-decoder architecture with two unified conditional encoders. By alternatively switching the two conditional encoders, we are able to train the NMT predictor on both the primal task of source-to-target translation and the dual task of target-to-source translation, with double the training data compared to the training on a single task. Fig.2

gives an illustration of the model-level dual learning of the NMT predictor consisting of two conditional encoders. The detailed training procedures of the NMT predictor is demonstrated in algorithm 1, where the model is trained alternatively on the primal and dual translation tasks.

**3.1.2 Mixture Models for MT.** The one-to-many mapping property of machine translation is one critical reason why it is a challenging task, as multiple valid translations exist for a given source sentence. The auto-regressive setting of transformer makes it difficult to search for multiple modes of solution methods, due to the fact that all uncertainty is represented in the decoder output distribution.

Commonly used decoding algorithms such as beam search allow some diversity in the decoding results. But they are limited to produce only minor differences in the suffixes. Mixture models are another option to model uncertainty and generate diverse translations through different mixture components.

Formally,  $K$  expert mixture models are trained and combined, where in each mixture model, a multinomial latent variable  $z \in \{1, \dots, K\}$  is introduced and the objective of machine translation is decomposed as

$$p(y|x) = \sum_{z=1}^K p(y, z|x) = \sum_{z=1}^K p(z|x)p(y|z, x), \quad (2)$$

where the prior  $p(z|x)$  and likelihood  $p(y|z, x)$  are learned functions.

The posterior probability  $p(z|x, y)$  can be utilized to determine the responsibility of each expert, and weigh the gradients of parameter update for each expert. However, in this paper, we use a hard assignment of responsibilities, where each expert  $z'$  is used to generate a hypotheses, and the expert giving the highest joint probability takes full responsibilities while others take none:

$$r_z = \begin{cases} 1, & z = \arg\max_{z'} p(y, z'|x), \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

This hard assignment encourages the implication that each particular sentence should be handled only by one expert who is specialized on it, so that each expert mixture model can individually master its own translation styles.

Moreover, we set the prior distribution  $p(z|x)$  to uniform in order to push all experts to generate good hypotheses for any sample and alleviate the degeneracy that only one expert gets trained well due to the “rich get richer” effect. The expert models are trained using the EM algorithm [18] and the objective function with uniform prior and hard assignments is formulated as:

$$L = \mathbb{E}_{(x,y) \sim D} \left[ \min_z -\log p(y|z, x) \right]. \quad (4)$$

For the design of the expert mixture models, there are two options of either utilizing an independent decoder for each expert or using a same decoder network but with different start-of-sentence token embeddings. We pick the latter option because it does not overly increase the model size and the shared decoder allows even the poor experts to be trained which further mitigates the degeneracy.

### 3.2 XLM Predictor

One of the hottest topic in nature language processing is the language model pre-training, where a language model is firstly trained on a large text corpus, and fine-tuned on specific NLP tasks. The pre-trained language models (PLMs) have achieved great success and become a must-have module in many NLP applications. Despite the effectiveness, most PLMs are trained on monolingual corpora and only focus on English language tasks, which can hardly fulfill the needs of multi-lingual tasks such as machine translation.

Cross-lingual Language Models (XLMs) [4] aim to build a universal cross-lingual language encoder that maps sentences of different languages into a shared embedding space. The pre-training of an XLM model involves both monolingual and cross-lingual objectives, including masked language modeling and translation language modeling.

Masked Language Modeling (MLM) is a typical objective in language model pre-training with unsupervised monolingual corpora, in which a portion of the input tokens are randomly selected, and replaced by [MASK] tokens, or random tokens, or stay unchanged. The altered input sentences are then passed through a transformer encoder to model the correlations between the input tokens, with the objective being to correctly recover the masked tokens based on their preceding and following contexts.

In addition to MLM, XLM models are also pre-trained on supervised cross-lingual corpora through Translation Language Modeling (TLM). The objective of TLM is an extension of MLM, that is also to recover the masked tokens. However, the inputs of TLM includes a concatenated translation sentence pair. And the words in both the source sentence and the target sentence are randomly masked. When predicting a masked word, the model can attend not only to surrounding words in the sentence, but also the words in the other language, encouraging the model to align the representations of the two languages.

Through the pre-training of MLM and TLM, given a translation pair, cross-lingual language models can effectively model the cross-lingual representations and produce meaningful features for quality estimation.

### 3.3 QE Feature Extraction

The predictors trained on the parallel corpus will serve as feature extractors for the quality estimation task. It should be noted that there exists discrepancy between the machine translation task and the quality estimation task due to the fact that unlike translation datasets, the target sentences in QE datasets are not the golden references with few translation errors. The mistakes contained in translated sentences are necessary for the learning of the quality estimation models to distinguish between good and bad translations. Due to this discrepancy, directly using the hidden representations outputted from the trained predictors would not be the optimal solution. In our design, we follow the idea of previous work [6] to include mis-matching features and model derived features, and propose a new design of dual model features based on the model-level dual learning.

**3.3.1 Mis-matching Feature.** The basic intuition of mis-matching features is to measure the difference between the given translation sentences and the prior knowledge of the predictors. Specifically,

for each position in the target sentence, the predictors will predict the probability  $p(y_i|x)$  for each word  $y_i$  in the target vocabulary. Theoretically, a well-trained predictor is leaning to output high probabilities for correct translations in the sentence, and output low probabilities for mis-translated or missing words.

Formally, we include four dimensions in the mis-translated feature vector for each token:

$$f_{mm} = [\log p(y_{mt}|x), \log p(y_{max}|x), \log p(y_{mt}|x) - \log p(y_{max}|x), \mathbf{1}_{y_{mt}=y_{max}}], \quad (5)$$

where  $\log p(\cdot)$  denotes the log probability function,  $y_{mt}$  is the token in the given translation, and  $y_{max}$  is the token with maximum probability output from the predictor.

**3.3.2 Model derived Feature.** Similar to QEFVs, in model derived features, the token embedding is often used in quality estimation models [6, 12] to project the hidden representation  $z$  derived from the predictors for each target token  $y$  to the desired representations for QE task. In our model, the hidden representation  $z$  is the decoder output before the final prediction layer, derived by first encoding the source sentence  $x$  with a conditional encoder  $EncA$  and decoding the encoder output  $h_x$  together with the target  $y$  with  $EncB$ ,

$$\begin{aligned} h_x &= EncA.encode(x) \\ z &= EncB.decode(h_x, y) \end{aligned} \quad (6)$$

The final feature  $f_{model} = z \odot e$  is derived as the element-wise product of the hidden representation  $z$  and the target embedding  $e$ .

**3.3.3 Dual Model Feature.** Model-level dual learning allows the NMT predictor to be trained on both translation directions, thereby, offering the possibility of solely encoding the target information without considering the source contexts. Based on the encoding, we propose a novel feature named Dual Model Feature.

Normally, in an NMT model, the target sentence information is only involved in the decoding stage with the source sentence encoding, to predict the next word in the target sentence while the following words are masked. Such training setting prevents us from retrieving a good representation of the target sentence. Fortunately, our NMT predictor is composed of conditional encoders and trained on the target-to-source translation task, meaning that the encoding of target sentences are also well trained. The encoding of the target sentence  $y$  is derived as

$$z' = EncB.encode(y). \quad (7)$$

And the dual model feature  $f_{dual} = z'_i \odot e_i$  is derived similarly to  $f_{model}$ .

Compared with the model derived feature  $f_{model}$  which extracts information based on source context and masked translation, our dual model feature  $f_{dual}$  focuses on the internal correlations between translation words and constructs a compact representation on the target side.

### 3.4 Estimator

After we extract the features from the predictors, an estimator is applied to learn a transformation function that maps the features to the estimation output. The task of word and sentence-level quality estimation can be treated as sequential labeling and regression problems. Therefore, we have designed our estimator based on

**Table 2: Description of the WMT20 QE dataset.**

	Overall			Sentence-Level			Word-Level		
Dataset	Num Samples	Avg Src-Len	Avg Tgt-Len	Avg HTER	Min HTER	Max HTER	Avg BAD Ratio	Min BAD Ratio	Max BAD Ratio
train	7,000	16.51	17.15	0.628	0.0	1.0	0.305	0.0	0.571
dev	1,000	16.31	17.06	0.603	0.0	1.0	0.290	0.0	0.526
test	1,000	16.77	17.23	–	–	–	–	–	–

\* We report the statistics of the WMT20 QE dataset including the average sentence length of the source and target sentences, and the average, minimum, and maximum value of HTER score and the proportion of BAD tokens in all tokens.

a bidirectional GRU network connected with a few feed-forward layers.

The designed features are first concatenated together along the depth direction for both the NMT predictor and the XLM predictor. The concatenated features are then fed to the a bi-GRU network followed by a few feed-forward layers to predict the HTER score or word-level tags for the sentence and word-level QE respectively. In order to alleviate the overfitting problem, the weights of the Bi-GRU networks handling NMT and XLM features are shared.

## 4 EXPERIMENTS

Our proposed framework is evaluated on the WMT20 English-Chinese word and sentence-level quality estimation tasks. We have also conducted experiments on filtering bilingual corpus based on the evaluation results to show the potential application of our proposed QE model. We will first introduce the dataset used in our evaluation and provide the implementation details along with the performance results.

### 4.1 Datasets, Metrics and Baselines

In our *Verdi* framework, a parallel translation dataset is used to train the NMT predictor, and the estimator is trained on the QE dataset. For the machine translation task, we choose the CCMT-20 Chinese-English translation dataset and filter out samples with sentence length greater than 70, resulting in 8.95M data samples. The average lengths for English and Chinese sentences are 19.66 and 18.61 respectively.

The QE dataset contains 7,000 training samples, 1,000 validation samples, and 1,000 on-line test samples. Ground-truth labels are not accessible for the on-line test samples, and the evaluation is performed on-line by submitting predictions to the on-line system. For each sample, a source sentence and its translation output from a machine translation system are present. And the labels include a sentence-level HTER score defined as the ratio of the number of edits required on the translation to its sentence length, and word-level tags for each token and gap indicating whether a token is correctly translated and if there is missing tokens in the translation sentence.

The detailed statistics of the QE dataset is summarized in table 2. We can see that the overall translation quality is not satisfying with an average HTER score above 0.6 and the BAD tokens proportion around 30%. Among different samples, the translation quality varies a lot. Perfect translations are present in the dataset with 0.0 HTER and no BAD tokens, whereas poorly translated sentences of 1.0 HTER and with 50% BAD tokens also exist.

We report the evaluation results with different metrics for word and sentence-level QE tasks shown as follows. Note that the evaluation of word-level QE is done with word tags and gap tags combined together, as requested in the WMT20 QE competition.

- **Pearson Correlation Coefficient** [2] is a statistics that measures linear correlation between two variables. It has a value between  $-1$  and  $+1$ , where a value of 0 means no linear correlation, and values greater than 0 indicates positive correlations and vise versa. It is reported for the sentence-level QE task.
- **MAE and RMSE**, i.e., mean absolute error and root mean squared error, are two of the most common metrics used to measure accuracy of continuous variables. They are also reported for the sentence-level QE task.
- **Matthews Correlation Coefficient (MCC)** [17] is used to measure the quality of binary classifications. Its value also ranges from  $-1$  to  $+1$ , where  $-1$  and  $+1$  mean total disagreement and total agreement, and 0 is close to random guesses. The MCC metric is reported in the word-level QE task.
- **F1-BAD and F1-OK** are the two F1 scores for tokens labeled with “BAD” and “OK”, where the F1 score is a measure of accuracy calculated from the precision and recall values.

In the experiments, we compare our methods with the following baseline methods.

- **NUQE** [16] uses a lookup layer containing embeddings for target words and their source-aligned words. The embeddings are concatenated and fed into two sets of feed-forward layers and a Bi-GRU layer. NUQE does not require training on parallel corpora.
- **PredEst** [12] is a OpenKiwi implementation [11] of the predictor-estimator framework, which uses a bidirectional LSTM and two unidirectional LSTMs as the predictor, and the estimator takes features produced by the predictor and makes the predictions.
- **XLM-PredEst and Bert-PredEst** [10] combine pre-trained language models with the predictor-estimator framework through transfer learning approaches. Pre-trained XLM and Bert models are utilized respectively in these two baseline methods.
- **Bilingual Expert** [6] is a QE system proposed by Fan et al.. They build a conditional target language model with a bidirectional transformer and propose to extract mis-matching features from the predictor. A Bi-LSTM model is used as the estimator.

**Table 3: Evaluation results of different models on WMT20 sentence-level and word-level QE tasks.**

Model	Sentence-Level			Word-Level		
	Pearson	MAE	RMSE	MCC	F1-BAD	F1-OK
NUQE [16]	0.4661	0.1513	0.1855	0.4900	0.6492	0.7826
PredEst [12]	0.5058	0.1470	0.1814	0.5091	0.6576	0.8491
Bilingual Expert [6]	0.5463	0.1429	0.1762	0.5427	0.6616	0.8768
Bert-PredEst [10]	0.5676	0.1449	0.1790	0.5501	0.6883	0.8226
XLM-PredEst [10]	0.5991	0.1382	0.1713	0.5627	0.6972	0.8577
<i>Verdi</i>	<b>0.6353</b>	<b>0.1342</b>	<b>0.1665</b>	<b>0.5806</b>	<b>0.7021</b>	<b>0.8785</b>
WMT20 winner (ensemble)	0.6641	0.1293	0.1596	0.5872	<b>0.7137</b>	0.8662
<i>Verdi</i> (ensemble)	<b>0.6672</b>	<b>0.1280</b>	<b>0.1588</b>	<b>0.5889</b>	0.7006	<b>0.8865</b>

\* We compare our single model of *Verdi* with other single model baselines and our ensemble model with the winner of WMT20 QE competition.

## 4.2 Experimental Settings

Our implementation is based on the Fairseq sequence modeling toolkit [19]. And we use four Tesla P40 gpus to train the NMT predictor and one Tesla P40 to train the estimator.

Before we train the NMT predictor, we perform preprocessing on the CCMT20 English-Chinese translation dataset. We first use LTP [3] to tokenize the Chinese sentences and perform BPE tokenization on both English and Chinese languages. And we construct a shared vocabulary of size 70,000 containing both English and Chinese tokens.

We construct our NMT predictor based on the transformer-base architecture defined in [27], with 6 encoding and decoding layers, 8 attention heads, 2,048 units for the feed-forward layers, and an embedding dimension of 512. The mixture model technique is applied to the NMT predictor by adding 5 different start-of-sentence tokens to the vocabulary as 5 experts. For the XLM predictor, we utilized an official released model pre-trained on both Masked Language Model and Translation Language Modeling tasks.

During the feature extraction from the NMT predictor, we extract from all the 5 expert mixture models and output a feature vector with  $5 \times 512$  dimensions for both model-derived features  $f_{model}$  and dual model features  $f_{dual}$ . Together with the 4-dimension mismatching features, the final QE feature for each token has a dimension of 5140. As dual model features cannot be derived from the XLM predictor, we replicate the model derived features to fill in the position.

The estimator consists of a bidirectional GRU with a hidden size of 512, followed by a few layers that differs in the word and sentence-level QE tasks. In sentence-level QE, we perform a top-k operation with  $k = 3$  followed by max-pooling on the Bi-GRU outputs of both the NMT features and the XLM features. The outputs are concatenated and fed into two linear layers to make the final prediction. MSE loss is used to update the network parameters.

In word-level QE, we first integrate the features of BPE tokens that belong to the same word into one compact representation by mean-pooling. The processed features are then handled by the Bi-GRU network. And different from sentence-level QE, we pass the Bi-GRU outputs directly to two FC layers with no pooling operations. In practice, we discover that the convergence on NMT

and XLM features is hard to synchronize. Therefore, we propose to use a weighted average with weights of 0.8 and 0.2 to combine the NMT and XLM predictions. Cross entropy loss is utilized to update the network. We treat gap tag prediction as a separate task and apply the same method as predicting word tags.

As to the ensemble method, we utilize the stacking architectures [29] where the outputs of various single systems are used as inputs in a second-phase regression algorithm. To avoid overfitting, we use a five-fold cross validation, as described in [16].

## 4.3 Main Results

Table 3 first compares the performance of our proposed single model with the mentioned baseline methods, and then compares our ensemble model with the winner of WMT20 QE competition on sentence and word levels, in terms of different evaluation metrics. We can see that our single method achieves the best performance on both word and sentence-level QE tasks and significantly outperforms other baseline algorithms. On the sentence-level QE dataset, given source sentences and their corresponding translations, the Pearson correlation coefficient, MAE, and RMSE of our single model result are 0.635, 0.134, and 0.167 respectively, while the best result of previous methods are 0.599, 0.138, and 0.171 respectively. Similarly, our method also achieves the best results among all methods in word-level QE task. We also compare our ensemble result to the winning system of WMT20 QE competition and show that our ensemble model is able to beat them in both word and sentence-level quality estimation.

The reason of our success is attributed to the different strategies applied in our framework to enhance the feature extraction ability. First, we utilized an XLM predictor together with an NMT predictor to improve the feature representations. Second, we use model-level dual learning and mixture model technique to train the NMT predictor to produce more diverse and robust QE features. Third, taking the advantage of model-level dual learning, we propose a novel dual model feature to model the internal correlations and construct a compact representation on the translated sentence, which further improves the QE features. By incorporating these strategies and modules, our *Verdi* framework is able to outperform other method by a significant margin.



**Table 4: Ablation test results of different model variants on WMT20 sentence-level and word-level QE task.**

Model	Sentence-Level			Word-Level		
	Pearson	MAE	RMSE	MCC	F1 BAD	F1 OK
NMT only	0.5743	0.1396	0.1725	0.5651	0.6847	<b>0.8792</b>
+ XLM predictor	0.6090	0.1375	0.1688	0.5663	0.6893	0.8768
+ mixture & dual learning	0.6278	0.1345	<b>0.1662</b>	0.5789	0.7006	0.8783
+ dual model feature $f_{dual}$	<b>0.6353</b>	<b>0.1342</b>	0.1665	<b>0.5806</b>	<b>0.7021</b>	0.8785

**Table 5: Evaluation results of the bilingual corpus filtering on CCMT20 English-Chinese dataset.**

Dataset	Num Samples	BLEU
Unfiltered	6, 000, 000	20.39
QE filtered	5, 000, 000	<b>20.48</b>

\* We report the average of the BLEU scores measured on the best three checkpoints.

#### 4.4 Analysis

We evaluate the impact of different modules and methods by performing an ablation test for both word and sentence-level QE on WMT20 datasets. Table 4 lists the performances of our model variants with different modules and methods included.

We incrementally accommodate different modules and methods to a base model which is a simple model using predictor-estimator architecture with a single transformer NMT predictor, with mismatching features and model derived features incorporated. The base model produces a result of 0.574 Pearson on sentence-level QE and 0.565 MCC on word-level QE. Based on it, we include the second XLM predictor and achieves 0.609 and 0.566 on sentence and word level respectively, which evidently proves the effectiveness of utilizing multiple predictors as feature extractors. Furthermore, we apply model-level dual learning and mixture model technique to the NMT predictor. This model variant further produces a performance of 0.628 Pearson and 0.579 MCC on the two tasks, and the significant improvement can be attributed to the model’s capability of generating better representations by exploiting the translation dualities and producing diverse translation outputs. Finally, we incorporate our designed dual model features to the model, and achieve the best performance of 0.635 Pearson score and 0.581 MCC score, showing that constructing a more compact representation of the translation sentence is helpful to the model performance.

#### 4.5 Applied to Bilingual Corpus Filtering

We also perform an experiment on bilingual corpus filtering by applying our QE model to evaluate the quality of samples in the parallel corpus. Specifically, we run the sentence-level evaluation on 6 million samples from the CCMT20 Chinese-English dataset with a trained QE model of our framework. We sort the 6 million samples according to their predicted HTER scores, and filter out 1 million samples with the highest error rates, resulting in a QE filtered dataset with 5 million translation pairs.

Table 5 compares the performance of transformer NMT model trained on the unfiltered 6 million samples and the filtered 5 million samples. We can see that even though the filtered dataset contains fewer samples, the model trained on it produces a better performance with a BLEU score of 20.48 compared to 20.39 given by the model trained on the full 6 million data. This result proves that the QE evaluation result is a good indication of the translation quality and can be potentially helpful to bilingual corpus filtering. By removing the low quality samples, we are able to reduce the training efforts, and at the same time, improve the performance of the NMT system, which from another aspect proves the effectiveness of our QE framework.

### 5 RELATED WORK

In this section, we review related work on quality estimation and dual learning.

The traditional studies on QE are mostly based on manual feature engineering that investigates the useful features to learn a shallow regressor/classifier to evaluate the quality of a given translation pair. Different features have been investigated including linguistic features [7], statistical baseline features [25], topic model features [21], and pseudo-reference and back-translation features [13], possibly with feature selection strategies such as principal component analysis [8]. However, the manually designed features require expert knowledge and are hard to transfer between datasets and models.

Deep learning has also been applied to quality estimation. Context window-based approaches are widely used in many deep learning QE studies, where bilingual context windows are derived for both source and target sentences with the center words in both languages properly aligned. Kreutzer et al. [14] proposes to adopt an FNN architecture on the context windows and combine it with a linear model using handcrafted QE features. Recurrent networks and convolutional networks on context windows are also introduced [15, 16, 20].

Recently, the predictor-estimator framework [12] is proposed to evaluate the translation quality without the need of word alignments. It consists of a word prediction RNN trained on parallel corpora that extracts representative features from the source sentence and its translation, and a neural quality estimation model that makes predictions based on the extracted features. Fan et al. [6] proposes to use a bidirectional transformer language model to produce the joint representation and designs features that measure the degree of mis-matching. Pre-trained language models such as BERT [5], are also introduced into the predictor-estimator framework as the feature extractor.



Different from the existing QE methods, our work utilizes an transformer predictor together with an XLM predictor, and incorporates model-level dual learning to enhance the feature extraction ability.

Dual learning is a technique that leverages the symmetric structure of some learning problems such as machine translation. He et al. [9] focuses on the unsupervised setting and proposes to assist the learning of a translation model by exploiting the distortion between a sentence and its back-translation. Dual supervised learning [30] conducts joint learning from labeled data by adding an additional probabilistic constraint. Model-level dual learning [31], from another aspect, exploits the duality of tasks in model level. It designs an architecture for primal and dual tasks, and ties the model parameters playing similar roles.

## 6 CONCLUSION AND FUTURE WORK

In this paper, we present our *Verdi* framework to solve the quality estimation problem. Our predictor module includes a pre-trained XLM model and an NMT model that is trained through model-level dual learning and mixture models technique. Besides, we also designed a novel feature to construct a better representation on the target side. Experiments conducted on WMT20 QE competition demonstrate that our system yields a better performance than other baselines and beats the winning system in the competition. We have also shown a potential application of employing our quality estimation method on bilingual corpus filtering task and demonstrated its effectiveness.

For future work, we plan to investigate the possibility of extending our framework to automatic post-editing tasks, in which the model does not only detect the potential mistakes, but performs automatic corrections to improve the translation qualities.

## REFERENCES

- [1] Daniel Beck, Kashif Shah, Trevor Cohn, and Lucia Specia. 2013. SHEF-Lite: When less is more for translation quality estimation. In *Proceedings of the Eighth Workshop on Statistical Machine Translation*. 337–342.
- [2] Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*. Springer, 1–4.
- [3] Wanxiang Che, Yunlong Feng, Libo Qin, and Ting Liu. 2020. N-LTP: A Open-source Neural Chinese Language Technology Platform with Pretrained Models. *arXiv preprint arXiv:2009.11616* (2020).
- [4] Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*. 7059–7069.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [6] Kai Fan, Jiayi Wang, Bo Li, Fengming Zhou, Boxing Chen, and Luo Si. 2019. “Bilingual Expert” Can Find Translation Errors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6367–6374.
- [7] Mariano Felice and Lucia Specia. 2012. Linguistic features for quality estimation. In *Proceedings of the Seventh Workshop on Statistical Machine Translation*. 96–103.
- [8] Jesús González-Rubio, Alberto Sanchis, and Francisco Casacuberta. 2012. PRHLT submission to the WMT12 quality estimation task. In *Proceedings of the seventh workshop on statistical machine translation*. 104–108.
- [9] Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *Advances in neural information processing systems*. 820–828.
- [10] Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, António Góis, M Amin Farajian, António V Lopes, and André FT Martins. 2019. 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)*. 78–84.
- [11] Fabio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, and André FT 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*. 117–122.
- [12] Hyun Kim, Hun-Young Jung, Hongseok Kwon, Jong-Hyeok Lee, and Seung-Hoon Na. 2017. Predictor-estimator: Neural quality estimation based on target word prediction for machine translation. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)* 17, 1 (2017), 1–22.
- [13] Anna Kozlova, Mariya Shmatova, and Anton Frolov. 2016. Ysda participation in the wmt’16 quality estimation shared task. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*. 793–799.
- [14] 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*. 316–322.
- [15] André FT 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*. 806–811.
- [16] André FT Martins, Marcin Junczys-Dowmunt, Fabio N Kepler, Ramón Astudillo, Chris Hokamp, and Roman Grundkiewicz. 2017. Pushing the limits of translation quality estimation. *Transactions of the Association for Computational Linguistics* 5 (2017), 205–218.
- [17] Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure* 405, 2 (1975), 442–451.
- [18] Geoffrey J McLachlan and Thriyambakam Krishnan. 2007. *The EM algorithm and extensions*. Vol. 382. John Wiley & Sons.
- [19] 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 NAACL-HLT 2019: Demonstrations*.
- [20] Raj Nath Patel and M Sasikumar. 2016. Translation Quality Estimation using Recurrent Neural Network. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*. 819–824.
- [21] Raphael Rubino, Jose de Souza, Jennifer Foster, and Lucia Specia. 2013. Topic models for translation quality estimation for gisting purposes. (2013).
- [22] Kashif Shah, Trevor Cohn, and Lucia Specia. 2015. A bayesian non-linear method for feature selection in machine translation quality estimation. *Machine Translation* 29, 2 (2015), 101–125.
- [23] Tianxiao Shen, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. 2019. Mixture Models for Diverse Machine Translation: Tricks of the Trade. In *International Conference on Machine Learning*. 5719–5728.
- [24] Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of association for machine translation in the Americas*, Vol. 200. Cambridge, MA.
- [25] 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*. Association for Computational Linguistics, Sofia, Bulgaria, 79–84. <https://www.aclweb.org/anthology/P13-4014>
- [26] Nicola Ueffing and Hermann Ney. 2007. Word-level confidence estimation for machine translation. *Computational Linguistics* 33, 1 (2007), 9–40.
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [28] Wei Wang, Taro Watanabe, Macduff Hughes, Tetsuji Nakagawa, and Ciprian Chelba. 2018. Denoising Neural Machine Translation Training with Trusted Data and Online Data Selection. In *Proceedings of the Third Conference on Machine Translation: Research Papers*. 133–143.
- [29] David H Wolpert. 1992. Stacked generalization. *Neural networks* 5, 2 (1992), 241–259.
- [30] Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017. Dual supervised learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. 3789–3798.
- [31] Yingce Xia, Xu Tan, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2018. Model-level dual learning. In *International Conference on Machine Learning*. 5383–5392.