

Translation Error Detection as Rationale Extraction

Marina Fomicheva

University of Sheffield

m.fomicheva@sheffield.ac.uk

Lucia Specia

Imperial College London

l.specia@imperial.ac.uk

Nikolaos Aletras

University of Sheffield

n.aletras@sheffield.ac.uk

Abstract

Recent Quality Estimation (QE) models based on multilingual pre-trained representations have achieved very competitive results in predicting the overall quality of translated sentences. However, detecting specifically which translated words are incorrect is a more challenging task, especially when dealing with limited amounts of training data. We hypothesize that, not unlike humans, successful QE models rely on translation errors to predict overall sentence quality. By exploring a set of feature attribution methods that assign relevance scores to the inputs to explain model predictions, we study the behaviour of state-of-the-art sentence-level QE models and show that explanations (i.e. rationales) extracted from these models can indeed be used to detect translation errors. We therefore (i) introduce a novel semi-supervised method for word-level QE; and (ii) propose to use the QE task as a new benchmark for evaluating the plausibility of feature attribution, i.e. how interpretable model explanations are to humans.

1 Introduction

Quality Estimation (QE) is the task of predicting Machine Translation (MT) quality at inference time, when no gold standard human translation is available (Blatz et al., 2004; Specia et al., 2009). QE can be framed as a word-level or a sentence-level task. Both tasks have numerous practical applications, such as deciding whether a given MT output can be published without editing, highlighting potential critical errors. Current QE approaches fine-tune powerful representations from pre-trained multilingual encoders such as BERT (Devlin et al., 2018) or XLM-R (Conneau et al., 2019). In the recent Shared Task on QE at WMT2020 (Specia et al., 2020) these approaches have achieved very high performance at predicting sentence-level translation quality (up to 0.9 Pearson correlation with human judgements for some language pairs). How-

ever, as evidenced by these results, the accuracy of word-level prediction still leaves room for improvement. This is partly due to the limited amount of training data. Word-level error annotation is especially time-consuming and expensive, as it requires work from bilingual experts. In this work we introduce a new semi-supervised approach to word-level QE that removes the need of training data at word level. To achieve this, we propose addressing QE as a rationale extraction task (Lei et al., 2016).

Explainability is a broad area aimed at explaining predictions of machine learning models (Lipton, 2016). Rationale extraction methods achieve this by selecting a portion of the input that justifies model output for a given data point. In translation, human perception of quality is guided by the presence of translation errors (Freitag et al., 2021). We hypothesize that sentence-level QE models also rely on translation errors to make predictions. If that is the case, explanations for sentence-level predictions can be used to detect translation errors, thus removing the need for word-level labeled training data. To extract model explanations, we use *post hoc* rationale extraction methods (Sundararajan et al., 2017) which try to explain the predictions of a given model (as opposed to modifying its architecture or introducing constraints during training), since one of our goals is to study to what extent existing QE models rely on the same information as humans to make predictions.

At the same time, by using word-level errors as explanations for sentence-level QE scores, we introduce a new benchmark for evaluating explainability methods. Recent work has introduced various datasets for measuring the agreement between rationales extracted from NLP models and those provided by humans (DeYoung et al., 2019). QE is different from these datasets in various important aspects. First, it is a regression task, as opposed to binary or multiclass text classification mainly explored in previous work. Second, it is a multi-

lingual task where the output score captures the relationship between source and target sentences. Finally, manual annotation of translation errors is a practical task with a long tradition in MT research and translation studies (Lommel et al., 2014), and thus offers an interesting alternative to human explanations collected specifically for evaluating rationale extraction methods.

Our **main contributions** are:

- We introduce a novel semi-supervised approach for word-level QE. We provide practical recipes on how feature attribution methods can be used to derive information on translation errors from sentence-level models.
- We provide insights into the behaviour of state-of-the-art (SOTA) QE models by analysing attributions to different parts of the input sequence (source vs. target sentence, correct words vs. errors) at different hidden layers.
- We propose to use the QE task as a new benchmark for evaluating the plausibility of feature attribution explanations, i.e. how interpretable model explanations are to humans (Jacovi and Goldberg, 2020).

2 Background and Related Work

Quality Estimation Current SOTA models in *sentence-level* QE, which is typically framed as a regression task, mainly use multilingual representations from pre-trained transformers (Devlin et al., 2018), notably XLM-R (Conneau et al., 2019). The input to a sentence-level QE model is a concatenation of the source and translated sentences, separated by the [SEP] token. The sequence is encoded by the pre-trained Transformer model, and the [CLS] token is passed through a multilayer perceptron (MLP) layer to obtain a sentence-level score. During fine-tuning both the parameters of the pre-trained model and the parameters corresponding to the MLP layer are updated.

Word-level QE is typically addressed as a binary classification task, where the QE model needs to predict a binary label indicating whether a word is correct or wrong for each word in the MT output (Lee, 2020). As illustrated in Figure 1 (left), some supervised approaches use both sentence-level and word-level objectives in a multi-task setting, which results in superior performance (Kim et al., 2017;

Lee, 2020). Methods that do not require word-level training data either need access to the MT model (Riktors and Fishel, 2017; Fomicheva et al., 2020b), or still treat the problem as a supervised task but use synthetically generated data for supervision (Tuan et al., 2021).

Rationale Extraction for NLP SOTA NLP models based on deep neural networks achieve high performance in a variety of tasks, often at the cost of interpretability (Lipton, 2016). Recent work aims to address this issue by focusing on two different goals. On the one hand, the aim is to produce justifications for model predictions that are *plausible* to the users, in order to increase users’ trust (Ribeiro et al., 2016). On the other hand, the aim is to reveal the inner workings of the model and *faithfully* explain its predictions, so the explanation can be useful to model developers (Jacovi and Goldberg, 2020).

Typically, explainability methods operate by selecting a portion of the input that justifies model prediction for a single data point. This can be done either by modifying the model architecture, or by trying to explain the predictions of a given model. The first type of approaches (a.k.a. rationalization by construction) involves imposing restrictions on the generated rationales to satisfy certain constraints, e.g. compactness (Yu et al., 2019; Chalkidis et al., 2021). Note that such restrictions often result in lower performance and indeed are not guaranteed to explain the behaviour of an unconstrained model (Jain et al., 2020). The second type of approaches (the so called *post hoc*) usually rely on feature attribution methods, which assign an importance value to each input feature of a network (Sundararajan et al., 2017; Schulz et al., 2020). These methods do not allow for introducing useful biases during training, but focus on faithfully explaining model behaviour.

Feature attribution has a long tradition in image recognition tasks (Simonyan et al., 2013) and only recently have been applied to some NLP tasks, most commonly text classification (DeYoung et al., 2019). QE is fundamentally different from text classification where clues are typically separate words or phrases (Zaidan et al., 2007) which often can be considered independently of the rest of the text. This independence assumption does not hold for the task of evaluating translation quality where a word cannot be identified as a clue (e.g. translation error) without considering the surrounding context.

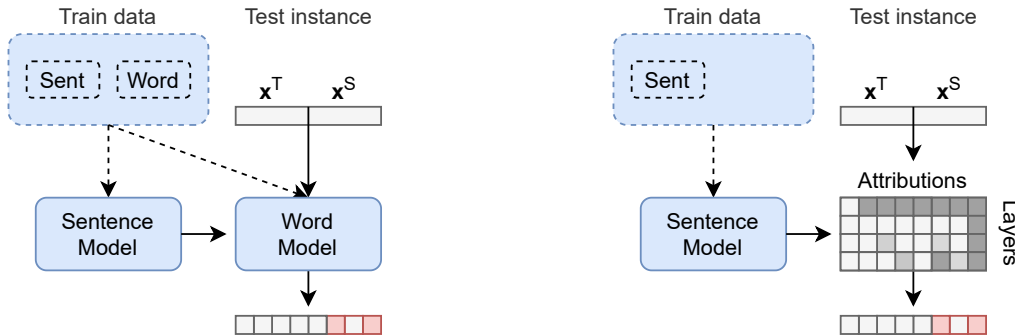


Figure 1: Fully supervised word-level QE (left) and semi-supervised word-level QE as rationale extraction (right). Dashed and solid lines represent training and test time, respectively.

Furthermore, SOTA NLP models based on contextualized representations for input words make rationale extraction especially challenging, as the representation for a given word can encode not only the word identity but also its interactions with other words in the text. Recent work has revealed various interesting properties that characterize the information flow through hidden layers in deep transformer models (Voita et al., 2019; De Cao et al., 2020; Yun et al., 2021). We provide additional insights on this topic in Section 5.2.

3 Translation Error Prediction as Rationale Extraction

We propose framing semi-supervised word-level QE as rationale extraction from sentence-level QE models. Instead of training a dedicated supervised model for word-level prediction, we propose deriving word-level scores from a strong sentence-level QE model by extracting explanations for model predictions (see Figure 1 (right)). Given a trained sentence-level QE model and the test data, rationale extraction methods detect the parts of the input that are relevant for model predictions on a sample-by-sample basis. We hypothesize that words with the highest relevance scores should correspond to actual translation errors on word-level.

3.1 Approach

More formally, given the source sequence $\mathbf{x}^S = x_1^S, \dots, x_{|S|}^S$, the target sequence $\mathbf{x}^T = x_1^T, \dots, x_{|T|}^T$ and the QE model $M(\mathbf{x}^S, \mathbf{x}^T) = \hat{y}$ that predicts sentence MT quality, a feature attribution method produces a vector of attribution scores $\mathbf{a} = a_1, \dots, a_{|S+T|}$, which represent the contribution of each source and target word to the prediction \hat{y} .

Crucially, no word-level labels are required for training. For evaluation, the attribution scores

are compared against binary gold labels $\mathbf{w} = w_1, \dots, w_{|T|} \in \{0, 1\}$ indicating whether each given word in the target sequence is an error or correct.

The predictive models for QE explored in our experiments are built by fine-tuning multilingual representations from pre-trained transformers. Transformer model starts from context-agnostic representations consisting of positional and token embeddings. These representations are passed through a set of hidden layers where at each layer the representations are iteratively updated via multi-head attention. This allows the hidden representation for each token to encode information on other words in the sentence.

We note that attribution to the input tokens or to the embedding layer can hardly succeed in detecting translation errors, as those cannot be identified independently from the context given by the source and target sentence. In this work, we perform feature attribution to hidden states at different layers and analyse which layer results in attribution scores that best correspond to translation errors.

3.2 Feature Attribution Methods

Feature attribution methods can be divided into those providing explanations by simplification, such as LIME (Ribeiro et al., 2016); gradient-based explanations (Sundararajan et al., 2017); and perturbation-based explanations (Schulz et al., 2020).

We select three popular methods for rationale extraction, which (i) do not require modifying the model architecture or re-training the model and (ii) allow attribution to hidden states. For comparison, we also use LIME which operates directly on the input text. We note that this set is not exhaustive of SOTA rationale extraction methods. Our main goal is not to conduct a comparative study of feature

attribution methods but rather testing whether it is possible to address word-level QE as a rationale extraction task without any word-level supervision.

LIME (Ribeiro et al., 2016) is a simplification-based explanation technique, which fits a sparse linear model in the vicinity of each test instance, to approximate the decision boundary of the complex model.¹ The data for fitting the linear model is produced by perturbing the given instance and computing model predictions. Linear model coefficients are then used as attribution scores for each input feature. For NLP tasks features correspond to input tokens and perturbation is achieved by randomly removing words from the sequence.

Information Bottleneck is a perturbation-based method originally proposed by Schulz et al. (2020) for the task of image recognition. The method applies the idea of information bottleneck (Tishby and Zaslavsky, 2015) for feature attribution. Specifically, it injects noise into an intermediate layer representation. The amount of noise injected at the position corresponding to each input feature is optimized to minimize the loss of the main task while at the same time maximizing the overall amount of injected noise.

Integrated Gradients (Sundararajan et al., 2017) is a gradient-based method similar to the traditional salience and input*gradients approaches (Simonyan et al., 2013). The latter takes the signed partial derivatives of the output with respect to the input and multiply them by the input itself. Intuitively, this is analogous to inspecting the products of model coefficients and feature values in linear models (Sundararajan et al., 2017). Integrated gradients improves on that by defining a baseline input and computing the average gradient while the input varies along a linear path from baseline input to the actual input. The baseline is defined by the user depending on the task. For image recognition, black image is used as baseline. It is not clear what such baseline representation should be in the case of language tasks. Here, we select a zero baseline for simplicity. Better results can be achieved with a more informed choice of a baseline and we leave this to future work.²

¹We use the implementation available at <https://github.com/marcotcr/lime>

²For both information bottleneck and integrated gradients method we adapt the implementation available at <https://github.com/nicola-decao/diffmask> for our QE scenario. Our code will be made available upon acceptance.

Attention Finally, we test attention as an attribution method. Self-attention mechanisms have been widely studied in the context of explainability (Jain and Wallace, 2019; Serrano and Smith, 2019; Bujel et al., 2021). To compute a single attention score for a transformer-based model with multi-head attention, we average the weights across the different attention heads.

4 Experimental Setup

4.1 Evaluation Metrics

Given a test set with both sentence-level and word-level gold labels, we want to measure to what extent the words with the highest attributions according to the QE model correspond to human annotations for MT errors. Note that we cannot use the evaluation metrics traditionally employed for assessing the performance of word-level QE, such as F1 score and Matthews correlation coefficient (Specia et al., 2020), as they require binary predictions while feature attribution methods return continuous scores. Instead, we rely on metrics based on class probabilities (Atanasova et al., 2020). Since attribution methods proceed on instance-by-instance basis and the scores produced for different instances are not necessarily comparable, we compute the evaluation metrics for each instance separately and average the results across all instances in the test set.

AUC score For each instance, we compute the area under the receiver operating characteristic curve (AUC score) to compare the continuous attribution scores \mathbf{a} against binary gold labels \mathbf{w} . For a test set with N instances:

$$AUC = \frac{1}{N} \sum_n AUC_n(\mathbf{w}_n, \mathbf{a}_n^T) \quad (1)$$

Average Precision AUC score can be overly optimistic for imbalanced data. Therefore, we also use Average Precision (AP).

Recall at Top-K In addition, we report the Recall-at-Top-K commonly used in information retrieval. Applied to our setting, this metric computes the proportion of words with the highest attribution that correspond to translation errors against the total number of errors in the MT output. Thus, for a given instance (we omit the instance index n here for simplicity):

$$\text{Rec@TopK} = \frac{1}{k} \sum_{j \in \mathbf{e}_{1:k}} \mathbf{w}_j \quad (2)$$

	Ro-En	Et-En	Ne-En
Pearson r	0.84	0.66	0.66
Average DA	68.9	55.2	36.6
Num. sentences (all data)	1,000	1,000	1,000
Num. sentences (DA < 70)	438	640	935
Error rate (all data)	0.21	0.28	0.65
Error rate (DA < 70)	0.35	0.36	0.66

Table 1: General statistics for MLQE-PE test sets: performance of sentence-level QE models (Pearson r), average DA score, total number of sentences in the test set, number of sentences with DA < 70, as well as error rate in the full test set and in the subset of selected sentences.

Where $\mathbf{e} = \text{argsort}(\mathbf{a}^{\mathbf{x}^T})$ is a sequence of indices corresponding to target words sorted by attribution score from highest to lowest and k is the number of errors in the sentence. We then average the result across all instances in the test set.

Accuracy at Top-1 Finally, we report the proportion of sentences where the word with the highest attribution in the target corresponds to a translation error.

$$\text{Acc@Top1} = \frac{1}{N} \sum I[\mathbf{a}_{\mathbf{e}_1} = 1] \quad (3)$$

We note that the above metrics are not defined for sentences where all words are labelled as errors or correct. We exclude such sentences from evaluation.

4.2 Sentence-level QE

For sentence-level QE, we rely on TransQuest (Ranasinghe et al., 2020b), which was one of the top submissions to the WMT20 QE Shared Task (Specia et al., 2020). To facilitate the use of feature attribution methods described above, we use our own implementation of the approach proposed by (Ranasinghe et al., 2020b,a). It achieves comparable results to the ones reported by the authors. Due to limited computational resources we use the XLM-R-base as the underlying pre-trained Transformer model. We expect that using a more powerful sentence-level model would result in higher performance.

4.3 Data

We use the MLQE-PE (Multilingual Quality Estimation and Post-Editing) dataset described in

Fomicheva et al. (2020a).³ MLQE-PE provides various types of manual MT evaluation for multiple language pairs. The MT outputs were assigned a sentence-level score inspired by the Direct Assessment (DA) annotation (Graham et al., 2015; Guzmán et al., 2019) on a continuous [0, 100] scale capturing overall translation quality. In addition, the MT outputs were independently post-edited by professional translators. MT outputs and their corresponding post-edited versions were automatically aligned in order to derive word-level binary labels (“BAD” if the word was corrected, and “OK” otherwise), as well as their HTER score that corresponds to the average number of “BAD” labels in a sentence (Snover et al., 2006). We use these labels to evaluate the performance of different feature attribution approaches. We treat “BAD” labels as the positive class and “OK” labels as negative class in all of our experiments.⁴ We do not evaluate attribution to source words.

It is worth noting that word-level labels derived from post-editing do not capture error severity and do not always correspond to translation errors. However, due to the costs of collecting detailed error annotations for the substantially large amounts of data required to train SOTA models, this is a standard way of approximating error annotation in QE (Specia et al., 2020).⁵

To circumvent the above limitation, we leverage both types of sentence-level annotation (DA and HTER scores) in our experiments. We train sentence-level QE models with (i) DA scores and (ii) HTER scores. We evaluate both types of models using the word-labels derived from post-editing as described above. We then conduct evaluation as follows:

1. We first evaluate explanations for DA-based models on the sentences with a sentence-level DA score lower than 70.⁶

³<https://github.com/sheffieldnlp/mlqe-pe>

⁴The tokenization used internally by XLM-R model is different from the tokenization used for producing word-level error labels. To map the attribution scores to the word labels we take their maximum value.

⁵Despite the limitations, we have chosen this dataset because it provides (i) sufficient amount of word-level training data, which allows us to compare our approach to a SOTA supervised approach; and (ii) access to the neural MT models that were used to produce the translations, thus enabling a comparison to an unsupervised glass-box approach.

⁶This threshold is selected based on the annotation guidelines described in Fomicheva et al. (2020a), as the sentences assigned a score lower than 70 are guaranteed to have translation errors.

Method	Romanian-English				Estonian-English				Nepalese-English			
	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K
Gradients	0.75	0.72	0.84	0.62	0.66	0.63	0.72	0.52	0.66	0.81	0.91	0.72
Info. Bottleneck	0.65	0.62	0.71	0.50	0.58	0.55	0.56	0.46	0.64	0.78	0.80	0.71
Attention	0.79	0.73	0.80	0.63	0.65	0.57	0.52	0.49	0.69	0.82	0.88	0.74
LIME	0.54	0.48	0.40	0.39	0.56	0.56	0.65	0.46	0.52	0.75	0.76	0.68
Random	0.50	0.43	0.36	0.33	0.50	0.47	0.38	0.37	0.50	0.70	0.62	0.65
Glass-box	0.74	0.66	0.66	0.55	0.69	0.63	0.65	0.54	0.64	0.79	0.78	0.73
MicroTransQuest	0.88	0.81	0.88	0.70	0.84	0.80	0.89	0.70	0.82	0.89	0.96	0.82

Table 2: AUC/AP scores, as well as accuracy at top-1 (A@1) and recall at top-K (R@K) for different rationale extraction methods on the test partition of MLQE-PE dataset. Best rationale extraction results are highlighted in bold. Attributions are computed with respect to the hidden states at layer 10.

2. We also evaluate explanations for DA-based sentence-level models on the full subset of sentences that contain at least one word-level error.
3. Finally, we evaluate explanations for HTER-based sentence-level models on the full subset of sentences that contain at least one word-level error.

Interestingly, despite the discrepancy between DA training objective and word labels derived from post-editing, explanations for DA-based models achieve better accuracy. We report the results for (1) in the main body of the paper, while (2) and (3) are reported in Appendix B.

We select three language pairs for our experiments: Estonian-English (Et-En), Romanian-English (Ro-En) and Nepali-English (Ne-En) with the best performance at sentence level achieved at WMT2020 Shared Task. Table 1 shows statistics for the respective test sets. These three language pairs present very different conditions for the task. Sentence-level model for Ro-En has much stronger performance in terms of Pearson correlation with human judgements. Ne-En has substantially lower translation quality where “BAD” words actually represent the majority class.

4.4 QE Benchmarks

We consider two benchmarks for word-level QE. On the one hand, we report the results for a strong supervised model based on pre-trained representations from XLM-R adapted to predict word-level binary labels derived from post-editing. To report the metrics presented in 4.1, we use the probability of the positive class as attribution scores. On the other hand, we consider a fully unsupervised

approach, which however, requires access to the neural MT model, that was used to generate the translations.

Black-box Supervised QE We use the word-level architecture available as part of the TransQuest toolkit (Ranasinghe et al., 2020b).⁷ Similarly to the sentence-level TransQuest model, it relies on XLM-Roberta-base pre-trained model fine-tuned for token classification task. We use XLM-Roberta-base to be consistent with the sentence-level settings.

Glass-box Unsupervised QE Fomicheva et al. (2020b) propose to extract information from the MT system to predict translation quality in a fully unsupervised way. Following their work, we use log-probabilities from the neural MT model as attribution scores. The lower the log-probability corresponding to each word, the higher the chance that this word constitutes an error.

5 Results

5.1 QE as Rationale Extraction

Table 2 shows the performance of our approach with different rationale extraction methods, as well as SOTA word-level QE methods for the MLQE-PE dataset. For the first three methods we compute the attributions to the hidden states at each layer on the dev set and report the results for this layer on the test set. First, our semi-supervised approach with all explanation methods substantially outperforms the random baseline.⁸ Among the different expla-

⁷https://tharindudr.github.io/TransQuest/architectures/word_level_architecture

⁸The smallest gap with respect to the random baseline is observed for Ne-En. The overall quality of the translation for this language pair is low. This setting might be less suitable

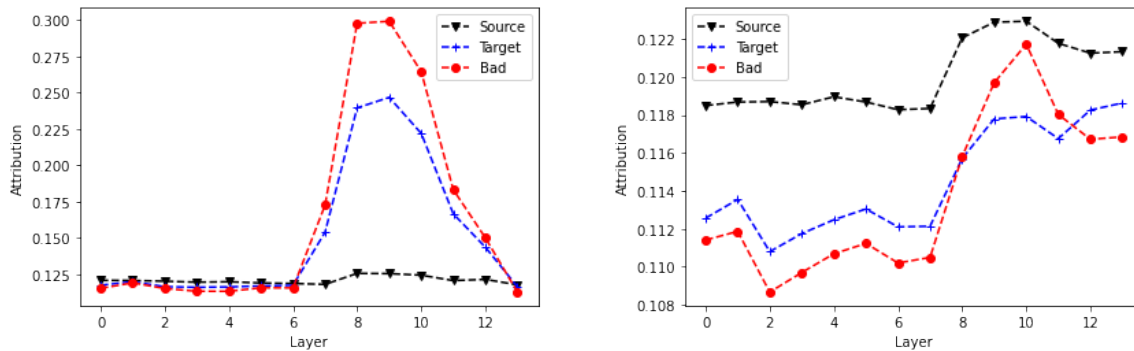


Figure 2: Average attribution at each hidden layer on the toy task (left) and MLQE-PE Et-En dataset (right). Attributions are computed with the information bottleneck attribution method (Schulz et al., 2020).

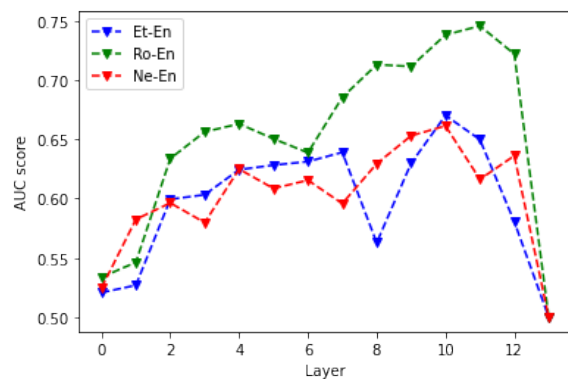


Figure 3: AUC score at each hidden layer for integrated gradients method.

nation methods, attention and integrated gradients achieve the best results. Second, the performance is comparable or better than the glass-box QE benchmark (MicroTransQuest) without requiring access to the neural MT model. For example, for Ro-En the AP scores achieved by the attention-based explanations and the glass-box word-level QE are 0.73 and 0.66, respectively. Third, the gap between the best-performing semi-supervised method and the supervised QE benchmark is the smallest for Ro-En, where the sentence-level QE model from which explanations are extracted is the strongest (see Table 1). Finally, on average, LIME-based explanations are substantially outperformed by the feature attribution methods. This agrees with our intuition that for the translation task where context plays a fundamental role, attribution to hidden states achieves much better performance than direct perturbation of input words.

for the proposed error detection methods as most of the words in the data correspond to errors, as shown in Table 1.

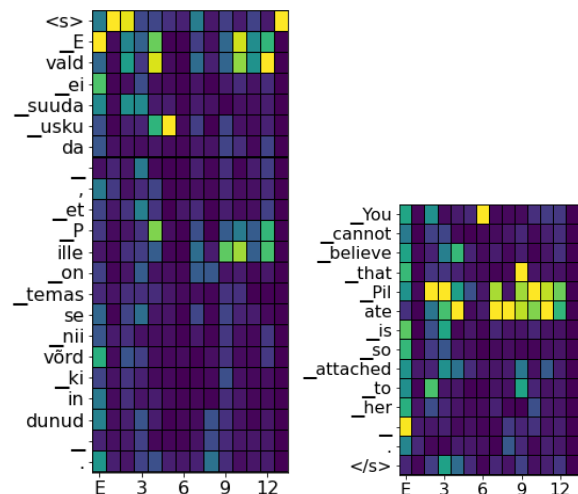


Figure 4: Example of Estonian-English translation with attributions to the source (left) and target (right) sentences computed using integrated gradients method for each hidden layer. The correct post-edited version of this translation is: *Evald cannot believe that Pille is so attached to her.*

5.2 Analysis

Feature Attribution per Layer Figure 2 shows attributions to tokens of different types across hidden layers. On the left, we show the results for a toy task, where we artificially introduced easy-to-detect errors in human translations and trained a QE model with near-perfect performance to predict whether a given sentence contain errors (see Appendix A). On the right, we show the results for the the MLQE-PE Et-En test set. Similarly to the toy task, we observe that in the later layers the tokens corresponding to translation errors receive higher attribution scores. However, in the toy dataset, the source tokens have very low attributions. Here, in

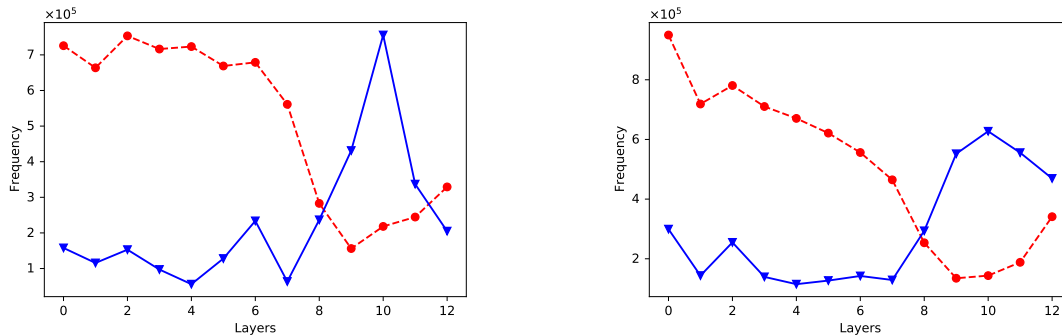


Figure 5: Frequency of the tokens with highest attribution in the neural MT training corpus. Y-axis shows the frequency of the source (left) and target (right) tokens with the highest attribution scores in low-quality MT sentences (red) and high-quality MT sentences (blue). X-axis corresponds to the hidden layers.

contrast, the model appears to be relying on the source as well as the target. This aligns very well with human evaluation where both source and target sentences need to be considered in order to correctly determine translation quality.

Figure 3 shows performance across layers for the integrated gradients method. As expected, the same layers that assign the highest attribution to the bad tokens (layers 9-11) are the ones that achieve the best performance. This finding is consistent across language pairs and attribution methods. Interestingly, this is also consistent with the findings reported in Voita et al. (2019), where they show that models trained with MLM objective encode context information in intermediate layers partially discarding the information on the identity of the input tokens which is recovered at the latest layers.

So far we have studied the behavior of the QE models on the sentences that contain errors. We now look at the pattern in the attributions scores for sentences which were assigned high quality by the model. We hypothesize that higher scores will be assigned to the words that are "easy" to translate. To test this, we select high-quality and low-quality sentences (sentences with predicted scores lower than 0.25 percentile and higher than 0.75 percentile, respectively). Figure 5 shows the average frequency with which the words occur in the neural MT training dataset. Red line corresponds to the words with the highest attribution for *high-quality* MT sentences. Blue line corresponds to the words with the highest attribution for the *low-quality* MT sentences. The first plot corresponds to the source tokens and the second plot corresponds to the target tokens. As shown in the plots, when the model predicts high quality the most frequent

words receive the highest attribution as the information progresses through the network. By contrast when low quality is predicted by the sentence-level model, the least frequent words receive the highest attribution.

Qualitative Analysis Figure 4 shows an example. Attributions are shown for sentencepiece tokens, which is the representation used internally by XLM-R. Interestingly, both translation errors ("You" and "Pilate") and the corresponding words in the source ("Evald" and "Pille") receive higher attribution scores.

6 Conclusion

In this work, we propose a new semi-supervised approach for word-level QE by exploring feature attribution methods. We show that for well performing models our results approach performance of supervised methods. We also consider the QE as rationale extraction task as a new benchmark for plausibility-based evaluation of explainability methods. We hope this work will encourage further research on improving the efficiency of word-level QE models with lightly supervised methods. This work opens many directions for future research: from improving the achieved results by tuning linear weights to combine attributions to hidden states at different layers, to exploring different underlying architectures and sentence-level training objectives.

Acknowledgements

This work was supported by funding from the Bergamot project (EU H2020 Grant No. 825303).

References

- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. A diagnostic study of explainability techniques for text classification. *arXiv preprint arXiv:2009.13295*.
- John Blatz, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto Sanchez, and Nicola Ueffing. 2004. [Confidence estimation for machine translation](#). In *Proceedings of the 20th International Conference on Computational Linguistics*, Geneva, Switzerland.
- Kamil Bujel, Helen Yannakoudakis, and Marek Rei. 2021. [Zero-shot sequence labeling for transformer-based sentence classifiers](#). In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP-2021)*, pages 195–205, Online. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapatsanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. Paragraph-level rationale extraction through regularization: A case study on european court of human rights cases. *arXiv preprint arXiv:2103.13084*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#). *arXiv preprint arXiv:1911.02116*.
- Nicola De Cao, Michael Schlichtkrull, Wilker Aziz, and Ivan Titov. 2020. How do decisions emerge across layers in neural models? interpretation with differentiable masking. *arXiv preprint arXiv:2004.14992*.
- 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*.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C Wallace. 2019. Eraser: A benchmark to evaluate rationalized nlp models. *arXiv preprint arXiv:1911.03429*.
- Marina Fomicheva, Shuo Sun, Erick Fonseca, Frédéric Blain, Vishrav Chaudhary, Francisco Guzmán, Nina Lopatina, Lucia Specia, and André FT Martins. 2020a. Mlqe-pe: A multilingual quality estimation and post-editing dataset. *arXiv preprint arXiv:2010.04480*.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020b. [Unsupervised quality estimation for neural machine translation](#). *arXiv preprint arXiv:2005.10608*.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. Experts, errors, and context: A large-scale study of human evaluation for machine translation. *arXiv preprint arXiv:2104.14478*.
- Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2015. [Can machine translation systems be evaluated by the crowd alone](#). *Natural Language Engineering*, pages 1–28.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. The flores evaluation datasets for low-resource machine translation: Nepali–english and sinhala–english. 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 6100–6113.
- Alon Jacovi and Yoav Goldberg. 2020. [Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness?](#) *CoRR*, abs/2004.03685.
- Sarthak Jain and Byron C Wallace. 2019. Attention is not explanation. *arXiv preprint arXiv:1902.10186*.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C Wallace. 2020. Learning to faithfully rationalize by construction. *arXiv preprint arXiv:2005.00115*.
- 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, Volume 2: Shared Tasks Papers*, pages 562–568, Copenhagen, Denmark.
- Dongjun Lee. 2020. [Two-phase cross-lingual language model fine-tuning for machine translation quality estimation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1024–1028, Online. Association for Computational Linguistics.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. [Rationalizing neural predictions](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas. Association for Computational Linguistics.
- Zachary Chase Lipton. 2016. [The myths of model interpretability](#). *CoRR*, abs/1606.03490.
- Arle Lommel, Hans Uszkoreit, and Aljoscha Burchardt. 2014. Multidimensional quality metrics (mqm): A framework for declaring and describing translation quality metrics. *Tradumática*, (12):0455–463.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020a. Transquest at wmt2020: Sentence-level direct assessment. *arXiv preprint arXiv:2010.05318*.

- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2020b. Transquest: Translation quality estimation with cross-lingual transformers. *arXiv preprint arXiv:2011.01536*.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Matiss Rikters and Mark Fishel. 2017. Confidence through attention. *arXiv preprint arXiv:1710.03743*.
- Karl Schulz, Leon Sixt, Federico Tombari, and Tim Landgraf. 2020. Restricting the flow: Information bottlenecks for attribution. *arXiv preprint arXiv:2001.00396*.
- Sofia Serrano and Noah A Smith. 2019. Is attention interpretable? *arXiv preprint arXiv:1906.03731*.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.
- 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*, volume 200. Citeseer.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. Findings of the WMT 2020 shared task on quality estimation. In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- Lucia Specia, Nicola Cancedda, Marc Dymetman, Marco Turchi, and Nello Cristianini. 2009. Estimating the sentence-level quality of machine translation systems. In *Proceedings of the 13th Annual Conference of the European Association for Machine Translation*, pages 28–35, Barcelona, Spain.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5. IEEE.
- Yi-Lin Tuan, Ahmed El-Kishky, Adithya Renduchintala, Vishrav Chaudhary, Francisco Guzmán, and Lucia Specia. 2021. Quality estimation without human-labeled data. *arXiv preprint arXiv:2102.04020*.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. The bottom-up evolution of representations in the transformer: A study with machine translation and language modeling objectives. *arXiv preprint arXiv:1909.01380*.
- Mo Yu, Shiyu Chang, Yang Zhang, and Tommi S Jaakkola. 2019. Rethinking cooperative rationalization: Introspective extraction and complement control. *arXiv preprint arXiv:1910.13294*.
- Zeyu Yun, Yubei Chen, Bruno A Olshausen, and Yann LeCun. 2021. Transformer visualization via dictionary learning: contextualized embedding as a linear superposition of transformer factors. *arXiv preprint arXiv:2103.15949*.
- Omar Zaidan, Jason Eisner, and Christine Piatko. 2007. Using “annotator rationales” to improve machine learning for text categorization. In *Human language technologies 2007: The conference of the North American chapter of the association for computational linguistics; proceedings of the main conference*, pages 260–267.

A Toy dataset

We devise a toy task to test feature attribution performance for word-level QE. We artificially introduce easy-to-detect errors in human translations and train a QE model with near-perfect performance to predict the presence/absence of such errors in a sentence. Specifically, we sample 10K/1K/1K sentence pairs from Es-En News-Commentary dataset (train/dev/test). Next, we artificially inject errors to half of the sentences at a rate of 0.1 using the following operations: insert, delete or replace random word, or swap two words selected at random.

We fine-tune an XLM-R-base model for a sentence-level binary classification task where sentences that contain errors are considered as positive class, and sentences that do not contain errors are considered as negative class. The F1-score of this sentence-level classifier is 0.97. This is expected as the task is very easy.

B Performance of Rationale Extraction Methods on HTER Data

Tables 4 and 5 show the performance of the proposed methods on the full subset of sentences that contain at least one word-level error for sentence-level QE models trained with HTER and DA ground truth scores. Pearson correlation for both types of models is shown in Table 3. Interestingly, even though for Ro-En and Et-En the performance of sentence-level models is near identical, extracted rationales are more accurate for the model trained with DA judgements.

	Ro-En	Et-En	Ne-En
Pearson r (DA)	0.84	0.66	0.66
Pearson r (HTER)	0.82	0.62	0.51
Num. sentences (all data)	1,000	1,000	1,000
Num. sentences (with errors)	714	889	945
Error rate (all data)	0.21	0.28	0.65
Error rate (with errors)	0.28	0.31	0.65

Table 3: Statistics for MLQE-PE test sets: performance of sentence-level QE models (Pearson r), total number of sentences with at least one translation error, and the error rate in the full test set and in the subset of sentences with at least one error.

Method	Romanian-English				Estonian-English				Nepalese-English			
	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K
Gradients	0.73	0.65	0.72	0.54	0.64	0.56	0.61	0.45	0.66	0.81	0.90	0.71
Info. Bottleneck	0.59	0.49	0.50	0.36	0.54	0.47	0.42	0.37	0.62	0.76	0.78	0.69
Attention	0.76	0.65	0.67	0.53	0.63	0.51	0.45	0.41	0.69	0.81	0.87	0.73
LIME	0.51	0.39	0.29	0.29	0.55	0.49	0.54	0.39	0.52	0.73	0.72	0.66
Random	0.50	0.38	0.27	0.25	0.50	0.41	0.34	0.31	0.50	0.70	0.63	0.64
Glassbox	0.73	0.59	0.55	0.48	0.70	0.58	0.59	0.48	0.64	0.78	0.77	0.72
MicroTransQuest	0.86	0.74	0.76	0.62	0.83	0.74	0.79	0.64	0.82	0.89	0.96	0.82

Table 4: AUC/AP scores, as well as accuracy at top-1 (A@1) and recall at top-K (R@K) for different rationale extraction methods on the MLQE-PE test set on the subset of sentences that contain at least one error for the sentence-level QE models trained to predict DA judgements.

Method	Romanian-English				Estonian-English				Nepalese-English			
	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K	AUC	AP	A@1	R@K
Gradients	0.69	0.59	0.61	0.48	0.66	0.59	0.66	0.49	0.64	0.77	0.82	0.70
Info. Bottleneck	0.53	0.43	0.38	0.32	0.58	0.50	0.47	0.38	0.57	0.73	0.68	0.67
Attention	0.74	0.61	0.59	0.49	0.69	0.59	0.58	0.48	0.66	0.78	0.82	0.72
LIME	0.61	0.47	0.37	0.35	0.64	0.56	0.59	0.45	0.53	0.74	0.76	0.68
Random	0.50	0.38	0.27	0.25	0.50	0.41	0.33	0.32	0.50	0.70	0.63	0.64
Glassbox	0.73	0.59	0.55	0.48	0.70	0.58	0.59	0.48	0.64	0.78	0.77	0.72
MicroTransQuest	0.86	0.74	0.76	0.62	0.83	0.74	0.79	0.64	0.82	0.89	0.96	0.82

Table 5: AUC/AP scores, as well as accuracy at top-1 (A@1) and recall at top-K (R@K) for different rationale extraction methods on the MLQE-PE test set on the subset of sentences that contain at least one error for the sentence-level QE models trained to predict HTER.