

Unbabel’s Participation in the WMT19 Translation Quality Estimation Shared Task

Fábio Kepler
Unbabel

Jonay Trénous
Unbabel

Marcos Treviso
Instituto de Telecomunicações

Miguel Vera
Unbabel

António Goís
Unbabel

M. Amin Farajian
Unbabel

António V. Lopes
Unbabel

André F. T. Martins
Unbabel

{kepler, sony, miguel.vera}@unbabel.com
{antonio.gois, amin, antonio.lopes, andre.martins}@unbabel.com
marcosvtreviso@gmail.com

Abstract

We present the contribution of the Unbabel team to the WMT 2019 Shared Task on Quality Estimation. We participated on the word, sentence, and document-level tracks, encompassing 3 language pairs: English-German, English-Russian, and English-French. Our submissions build upon the recent OpenKiwi framework: we combine linear, neural, and predictor-estimator systems with new transfer learning approaches using BERT and XLM pre-trained models. We compare systems individually and propose new ensemble techniques for word and sentence-level predictions. We also propose a simple technique for converting word labels into document-level predictions. Overall, our submitted systems achieve the best results on all tracks and language pairs by a considerable margin.

1 Introduction

Quality estimation (QE) is the task of evaluating a translation system’s quality without access to reference translations (Blatz et al., 2004; Specia et al., 2018). This paper describes the contribution of the Unbabel team to the Shared Task on Word, Sentence, and Document-Level (QE Tasks 1 and 2) at WMT 2019.

Our system adapts OpenKiwi,¹ a recently released open-source framework for QE that implements the best QE systems from WMT 2015-18 shared tasks (Martins et al., 2016, 2017; Kim et al., 2017; Wang et al., 2018), which we extend to leverage recently proposed pre-trained models

¹<https://unbabel.github.io/OpenKiwi>.

via transfer learning techniques. Overall, our main contributions are as follows:

- We extend OpenKiwi with a Transformer predictor-estimator model (Wang et al., 2018).
- We apply transfer learning techniques, fine-tuning BERT (Devlin et al., 2018) and XLM (Lample and Conneau, 2019) models in a predictor-estimator architecture.
- We incorporate predictions coming from the APE-BERT system described in Correia and Martins (2019), also used in this year’s Unbabel’s APE submission (Lopes et al., 2019).
- We propose new ensembling techniques for combining word-level and sentence-level predictions, which outperform previously used stacking approaches (Martins et al., 2016).
- We build upon our BERT-based predictor-estimator model to obtain document-level annotation and MQM predictions via a simple word-to-annotation conversion scheme.

Our submitted systems achieve the best results on all tracks and all language pairs by a considerable margin: on English-Russian (En-Ru), our sentence-level system achieves a Pearson score of 59.23% (+5.96% than the second best system), and on English-German (En-De), we achieve 57.18% (+2.44%).

2 Word and Sentence-Level Task

The goal of the word-level QE task is to assign quality labels (OK or BAD) to each *machine-translated word*, as well as to *gaps* between words

(to account for context that needs to be inserted), and *source words* (to denote words in the original sentence that have been mistranslated or omitted in the target). The goal of the Sentence-level QE task, on the other hand, is to predict the quality of the whole translated sentence, based on how many edit operations are required to fix it, in terms of HTER (Human Translation Error Rate) (Specia et al., 2018). We next describe the datasets, resources, and models that we used for these tasks.

2.1 Datasets and Resources

The data resources we use to train our systems are of three types: the QE shared task corpora, additional parallel corpora, and artificial triplets (*src*, *pe*, *mt*) in the style of the eSCAPE corpus (Negri et al., 2018).

- The En-De QE corpus provided by the shared task, consisting of 13,442 `train` triplets.
- The En-Ru QE corpus provided by the shared task, consisting of 15,089 `train` triplets.
- The En-De parallel dataset of 3,396,364 sentences from the IT domain provided by the shared task organizers. which we extend in the style of the eSCAPE corpus to contain artificial triplets. To do this, we use OpenNMT with 5-fold jackknifing (Klein et al., 2017) to obtain unbiased translations of the source sentences.
- The En-Ru eSCAPE corpus (Negri et al., 2018) consisting of 7,735,361 artificial triplets.

2.2 Linear Sequential Model

Our simplest baseline is the linear sequential model described by Martins et al. (2016, 2017). It is a discriminative feature-based sequential model (called LINEARQE). The system uses a first-order sequential model with unigram and bigram features, whose weights are learned by using the max-loss MIRA algorithm (Crammer et al., 2006). The features include information about the words, part-of-speech tags, and syntactic dependencies, obtained with TurboParser (Martins et al., 2013).

2.3 NuQE

We used NUQE (NeUral Quality Estimation) as implemented in OpenKiwi (Kepler et al., 2019) and adapted it to jointly learn MT tags, source tags and also sentence scores. We use the original architecture with the following additions. For learning sentence scores, we first take the average of

the MT tags output layer and then pass the result through a feed-forward layer that projects the result to a single unit. For jointly learning source tags, we take the source text embeddings, project them with a feed-forward layer, and then sum the MT tags output vectors that are aligned. The result is then passed through a feed-forward layer, a bi-GRU, two other feed-forward layers, and finally an output layer. The layer dimensions are the same as in the normal model. It is worth noting that NUQE is trained from scratch using only the shared task data, with no pre-trained components, besides Polyglot embeddings (Al-Rfou et al., 2013).

2.4 RNN-Based Predictor-Estimator

Our implementation of the RNN-based prediction estimator (PREDEST-RNN) is described in Kepler et al. (2019). It follows closely the architecture proposed by Kim et al. (2017), which consists of two modules:

- a *predictor*, which is trained to predict each token of the target sentence given the source and the left and right context of the target sentence;
- an *estimator*, which takes features produced by the *predictor* and uses them to classify each word as OK or BAD.

Our predictor uses a biLSTM to encode the source, and two unidirectional LSTMs processing the target in left-to-right (LSTM-L2R) and right-to-left (LSTM-R2L) order. For each target token t_i , the representations of its left and right context are concatenated and used as query to an attention module before a final softmax layer. It is trained on the large parallel corpora mentioned above. The estimator takes as input a sequence of features: for each target token t_i , the final layer before the softmax (before processing t_i), and the concatenation of the i -th hidden state of LSTM-L2R and LSTM-R2L (after processing t_i). We train this system with a multi-task architecture that allows us to predict sentence-level HTER scores. Overall, this system is capable to predict sentence-level scores and all word-level labels (for MT words, gaps, and source words)—the source word labels are produced by training a predictor in the reverse direction.

2.5 Transformer-Based Predictor-Estimator

In addition, we implemented a Transformer-based predictor-estimator (PREDEST-TRANS), follow-

ing Wang et al. (2018). This model has the following modifications in the *predictor*: (i) in order to encode the source sentence, the bidirectional LSTM is replaced by a Transformer encoder; (ii) the LSTM-L2R is replaced by a Transformer decoder with future-masked positions, and the LSTM-R2L is replaced by a Transformer decoder with past-masked positions. Additionally, the Transformer-based model produces the “mismatch features” proposed by Fan et al. (2018).

2.6 Transfer Learning and Fine-Tuning

Following the recent trend in the NLP community leveraging large-scale language model pre-training for a diverse set of downstream tasks, we used two pre-trained language models as feature extractors, the multilingual BERT (Devlin et al., 2018) and the Cross-lingual Language Model (XLM) (Lample and Conneau, 2019). The predictor-estimator model consists of a predictor that produces contextual token representations, and an estimator that turns these representations into predictions for both word level tags, and sentence level scores. As both of these models produce contextual representations for each token in a pair of sentences, we simply replace the predictor part by either BERT or XLM to create new QE models: PREDEST-BERT and PREDEST-XLM. The XLM model is particularly well suited to the task at hand, as its pre-training objective already contains a translation language modeling part.

For improved performance, we employ a pre-fine-tuning step by continuing their language model pre-training on data that is closer to the domain of the shared task. For the En-De pair we used the in-domain data provided by the shared task, and for the En-Ru pair we used the eSCAPE corpus (Negri et al., 2018).

Despite the shared multilingual vocabulary, BERT is originally a monolingual model, treating the input as either being from one language or another. We pass both sentences as input by concatenating them according to the template: [CLS] target [SEP] source [SEP], where [CLS] and [SEP] are special symbols from BERT, denoting beginning of sentence and sentence separators, respectively. In contrast, XLM is a multilingual model which receives two sentences from different languages as input. Thus, its usage is straightforward.

The output from BERT and XLM is split into

target features and source features, which in turn are passed to the regular *estimator*. They work with word pieces rather than tokens, so the model maps their output to tokens by selecting the first word piece of each token. For En-Ru the mapping is slightly different, it is done by taking the average of the word pieces of each token.

For PREDEST-BERT, we obtained the best results by ignoring features from the other language, that is, for predicting target and gap tags we ignored source features, and for predicting source tags we ignored target features. On the other hand, PREDEST-XLM predicts labels for target, gaps and source at the same time. As the predictor-estimator model, PREDEST-BERT and PREDEST-XLM are trained in a multi-task fashion, predicting sentence-level scores along with word-level labels.

2.7 APE-QE

In addition to traditional QE systems, we also use Automatic Post-Editing (APE) adapted for QE (APE-QE), following Martins et al. (2017). An APE system is trained on the human post-edits and its outputs are used as pseudo-post-editions to generate word-level quality labels and sentence-level scores in the same way that the original labels were created.

We use two variants of APE-QE:

- PSEUDO-APE, which trains a regular translation model and uses its output as a pseudo-reference.
- An adaptation of BERT to APE (APE-BERT) with an additional decoding constraint to reward or discourage words that do not exist in the source or MT.

PSEUDO-APE was trained using OpenNMT-py (Klein et al., 2017). For En-De, we used the IT domain corpus provided by the shared task, and for En-Ru we used the Russian eSCAPE corpus (Negri et al., 2018).

For APE-BERT, we follow the approach of Correia and Martins (2019), also used by Unbabel’s APE shared task system (Lopes et al., 2019), and adapt BERT to the APE task using the QE in-domain corpus and the shared task data as input, where the source and MT sentences are the encoder’s input and the post-edited sentence is the decoder’s output. In addition, we also employ a conservativeness penalty (Lopes et al., 2019), a

METHOD	TARGET F_1
STACKED LINEAR	43.88
POWELL	44.61

Table 1: Performance of the stacked linear ensemble and Powell’s method on the WMT17 dev set (F_1 -MULT on MT tags). The ensemble is over the same set of models³ reported in the release of the OpenKiwi (Kepler et al., 2019) framework. To estimate the performance of Powell’s method, the dev set was partitioned into 10 folds f_i . We ran Powell’s method 10 times, leaving out one fold at a time, to learn weights w_i . Predicting on fold f_i using weights w_i and calculating F_1 performance over the concatenation of these predictions gives an approximately unbiased estimate of the performance of the method.

beam decoding penalty which either rewards or penalizes choosing tokens not in the `src` and `mt`, with a negative score to encourage more edits of the MT.

2.8 System Ensembling

We ensembled the systems above to produce a single prediction, as described next.

Word-level ensembling. We compare two approaches:

- A stacked architecture with a feature-based linear system, as described by Martins et al. (2017). This approach uses the predictions of various systems as additional features in the linear system described in §2.2. To avoid overfitting on the training data, this approach requires jackknifing.
- A novel strategy consisting of learning a convex combination of system predictions, with the weights learned on the development set. We use Powell’s conjugate direction method (Powell, 1964)² as implemented in `SciPy` (Jones et al., 2001) to directly optimize for the task metric (F_1 -MULT).

Using the development set for learning carries a risk of overfitting; by using k -fold cross-validation we avoided this, and indeed the performance is equal or superior to the linear stacking ensemble (Table 1), while being computationally cheaper as only the development set is needed to learn an ensemble, avoiding jackknifing.

²This is the method underlying the popular MERT method (Och, 2003), widely used in the MT literature.

Sentence-level ensembling. We have systems outputting sentence-level predictions directly, and others outputting word-level probabilities that can be turned into sentence-level predictions by averaging them over a sentence, as in (Martins et al., 2017). To use all available features (sentence score, gap tag, MT tag and source tag predictions from all systems used in the word-level ensembles), we learn a linear combination of these features using ℓ_2 -regularized regression over the development set. We tune the regularization constant with k -fold cross-validation, and retrain on the full development set using the chosen value.

3 Document-Level Task

Estimating the quality of an entire document introduces additional challenges. The text may become too long to be processed at once by previously described methods, and longer-range dependencies may appear (e.g inconsistencies across sentences).

Both sub-tasks were addressed: estimating the MQM score of a document and identifying character-level annotations with corresponding severities. Note that, given the correct number of annotations in a document and their severities, the MQM score can be computed in closed form. However, preliminary experiments using the predicted annotations to compute MQM did not outperform the baseline, hence we opted for using independent systems for each of these sub-tasks.

3.1 Dataset

The data for this task consists of Amazon reviews translated from English to French using a neural MT system. Translations were manually annotated for errors, with each annotation associated to a severity tag (minor, major or critical).

Note that each annotation may include several words, which do not have to be contiguous. We refer to each contiguous block of characters in an annotation as a span, and refer to an annotation with at least two spans as a multi-span annotation. Figure 1 illustrates this, where a single annotation is comprised of the spans *bandes* and *parfaits*.

Across training set and last year’s development and test set, there are 36,242 annotations. Out of these, 4,170 are multi-span, and 149 of the multi-span annotations contain spans in different sentences. The distribution of severities is 84.12% of major, 11.74% of minor and 4.14% of critical.

Source: *resistance bands are great for home use, gym use, offices, and are ideal for travel.*

Target: *les **bandes** de résistance sont **parfaits** pour l'usage domestique, l'utilisation de la salle de gym, bureaux et sont idéales pour les voyages.*

Figure 1: Example of a multi-span annotation containing two spans: *parfaits* does not agree with *bandes* due to gender—it should be *parfaites*. This mistake corresponds to a single annotation with severity “minor”.

3.2 Implemented System

To predict annotations within a document the problem is first treated as a word-level task, with each sentence processed separately. To obtain gold labels, the training set is tokenized and an OK/BAD tag is attributed to each token, depending on whether the token contains characters belonging to an annotation. Note that besides token tags, we will also have gap tags in between tokens. A gap tag will only be labeled as BAD if a span begins and ends exactly in the borders of the gap. Our best-performing model for the word-level part is an ensemble of 5 BERT models. Each BERT model was trained as described in §2.6, but without pre-fine-tuning. Systems were ensembled by a simple average.

Later, annotations may be retrieved from the predicted word-level tags by concatenating contiguous BAD tokens into a single annotation. This is done for token-tags, while each gap-tag can be directly mapped to a single annotation without attempting any merge operation. Note that this immediately causes 4 types of information loss, which can be addressed in a second step:

- Severity information is lost, since all three severity labels are converted to BAD tags. As a baseline, all spans are assigned the most frequent severity, “major.”
- Span borders are defined on character-level, whose positions may not match exactly the beginning or ending of a token. This will cause all characters of a partially correct token to be annotated with an error.
- Contiguous BAD tokens will always be mapped to a single annotation, even if they belong to different ones.
- Non-contiguous BAD tokens will always be mapped to separate annotations, even if they belong to the same one.

PAIR	SYSTEM	TARGET F_1	SOURCE F_1	PEARSON
En-De	LINEAR	0.3346	0.2975	-
	APE-QE	0.3740	0.3446	0.3558
	APE-BERT	0.4244	0.4109	0.3816
	PRED-EST-RNN	0.3786	-	0.5020
	PRED-EST-TRANS	0.3980	-	0.5300
	PRED-EST-XLM	0.4144	0.3960	0.5810
	PRED-EST-BERT	0.3870	0.3310	0.5190
	LINEAR ENS.	0.4520	0.4116	-
	(*)POWELL'S ENS.	0.4872	0.4607	0.5968
En-Ru	LINEAR	0.2839	0.2466	-
	APE-QE	0.2592	0.2336	0.1921
	APE-BERT	0.2519	0.2283	0.1814
	NUQE	0.3130	0.2000	-
	PRED-EST-RNN	0.3201	-	-
	PRED-EST-TRANS	0.3414	-	0.3655
	PRED-EST-XLM	0.3799	0.3280	0.4983
	PRED-EST-BERT	0.3782	0.3039	0.5000
	(*)ENSEMBLE 1	0.3932	0.3640	0.5469
(*)ENSEMBLE 2	0.3972	0.3700	0.5423	

Table 2: Word and sentence-level results for En-De and En-Ru on the validation set in terms of F_1 -MULT and Pearson’s r correlation. (*) Lines with an asterisk use Powell’s method for word level ensembling and ℓ_2 -regularized regression for sentence level. As the weights are tuned on the dev set, their numbers can not be directly compared to the other models

Although more sophisticated approaches were tested for predicting severities and merging spans into the same annotation, these approaches did not result in significant gains, hence we opted by using the previously described pipeline as our final system. To predict document-level MQM, each sentence’s MQM is first predicted and used to get the average sentence MQM (weighting the average by sentence length degraded results in all experiments). This is used together with 3 percentages of BAD tags from the word-level model (considering token tags, gap tags and all gaps) as features for a linear regression which outputs the final document-level MQM prediction. The percentage of BAD tags is obtained from the previously described word-level predictions, whereas the sentence MQMs are obtained from an ensemble of 5 BERT models trained for sentence-level MQM prediction. Again, each BERT model was trained as described in §2.6 without pre-fine-tuning, and the ensembling consisted of a simple average.⁴

4 Experimental Results

4.1 Word and Sentence-Level Task

The results achieved by each of the systems described in §2 for En-De and En-Ru on the valida-

⁴Using the approach of [Ive et al. \(2018\)](#) proved less robust to this year’s data due to differences in the annotations. Particularly some outliers containing zero annotations would strongly harm the final Pearson score when mis-predicted.

PAIR	SYSTEM	TARGET F_1	TARGET MCC	SOURCE F_1	SOURCE MCC	PEARSON
En-Ru	<i>Baseline</i>	0.2412	0.2145	0.2647	0.1887	0.2601
	ENSEMBLE 1	0.4629	0.4412	0.4174	0.3729	0.5889
	ENSEMBLE 2	0.4780	0.4577	0.4541	0.4212	0.5923
En-De	<i>Baseline</i>	0.2974	0.2541	0.2908	0.2126	0.4001
	LINEAR ENSEMBLE	0.4621	0.4387	0.4284	0.3846	-
	POWELL’S ENSEMBLE	0.4752	0.4585	0.4455	0.4094	0.5718

Table 3: Word and sentence-level results for En-De and En-Ru on the test set in terms of F_1 -MULT and Pearson’s r correlation.

tion set are shown in Table 2. We tried the following strategies for ensembling:

- For En-De, we created a word-level ensemble system with Powell’s method, by combining one instance of the APE-BERT system, another instance of the PSEUDO-APE-QE system, 10 runs of the PREDEST-XLM model (trained jointly for all subtasks), 6 runs of the same model without pre-fine-tuning, 5 runs of the PREDEST-BERT model (trained jointly for all subtasks), and 5 runs of the PREDEST-TRANS model (trained jointly for MT and sentence subtasks, but not for predicting source tags). For comparison, we report also the performance of a stacked linear ensemble word-level system. For the sentence-level ensemble, we learned system weights by fitting a linear regressor to the sentence scores produced by all the above models.
- For En-Ru, we tried two versions of word-level ensemble systems, both using Powell’s method: EMSEMBLE 1 combined one instance of the APE-BERT system, 5 runs of the PREDEST-XLM model (trained jointly for all subtasks), one instance of the PREDEST-BERT model (trained jointly for all subtasks), 5 runs of the NUQE models (trained jointly for all subtasks), and 5 runs of the PREDEST-TRANS model (trained jointly for MT and sentence subtasks, but not for predicting source tags). EMSEMBLE 2 adds to the above predictions from the PSEUDO-APE-QE system. In both cases, for sentence-level ensembles, we learned system weights by fitting a linear regressor to the sentence scores produced by all the above models.

The results in Table 2 show that the transfer learning approach with BERT and XLM benefits the QE task. The PREDEST-XLM model, which has been pre-trained with a translation ob-

	DEV	DEV0	TEST
F_1 ANN. (BERT)	0.4664	0.4457	0.4811
MQM (BERT)	0.3924	-	0.3727
MQM (LINBERT)	-	0.4714	0.3744

Table 4: Results of document-level submissions, and their performance of the dev and dev0 validation sets.

jective, has a small but consistent advantage over both PREDEST-BERT and PREDEST-TRANSF. A clear takeaway is that ensembling of different systems can give large gains, even if some of the sub-systems are weak individually.

Table 3 shows the results obtained with our ensemble systems on the official test set.

4.2 Document-Level Task

Finally, Table 4 contains results for document-level submissions, both on validation and test set submissions. On F_1 annotations, results across all data sets are reasonably consistent. On the other hand, MQM Pearson varies significantly between dev and dev0. Differences in the training of the two systems shouldn’t explain this variation, since both have equivalent performance on the test set.

5 Conclusions

We presented Unbabel’s contribution to the WMT 2019 Shared Task on Quality Estimation. Our submissions are based on the OpenKiwi framework, to which we added new transfer learning approaches via BERT and XLM pre-trained models. We also proposed a new ensemble technique using Powell’s method that outperforms previous strategies, and we convert word labels into span annotations to obtain document-level predictions. Our submitted systems achieve the best results on all tracks and language pairs.

Acknowledgments

The authors would like to thank the support provided by the EU in the context of the PT2020 project (contracts 027767 and 038510), by the European Research Council (ERC StG DeepSPIN 758969), and by the Fundação para a Ciência e Tecnologia through contract UID/EEA/50008/2019.

References

- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed Word Representations for Multilingual NLP. *arXiv preprint arXiv:1307.1662*.
- John Blatz, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto Sanchis, and Nicola Ueffing. 2004. Confidence Estimation for Machine Translation. In *Proc. of the International Conference on Computational Linguistics*, page 315.
- Gonçalo Correia and André Martins. 2019. A Simple and Effective Approach to Automatic Post-Editing with Transfer Learning. In *Proc. of the 57th annual meeting on association for computational linguistics*. Association for Computational Linguistics.
- Koby Crammer, Ofer Dekel, Joseph Keshet, Shai Shalev-Shwartz, and Yoram Singer. 2006. Online Passive-Aggressive Algorithms. *Journal of Machine Learning Research*, 7:551–585.
- 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*.
- Kai Fan, Bo Li, Fengming Zhou, and Jiayi Wang. 2018. "Bilingual Expert" Can Find Translation Errors. *arXiv preprint arXiv:1807.09433*.
- Julia Ive, Carolina Scarton, Frédéric Blain, and Lucia Specia. 2018. Sheffield Submissions for the WMT18 Quality Estimation Shared Task. In *Proc. of the Third Conference on Machine Translation: Shared Task Papers*, pages 794–800.
- Eric Jones, Travis Oliphant, Pearu Peterson, et al. 2001. *SciPy: Open source scientific tools for Python*. [Online; accessed on May 17th, 2019].
- Fábio Kepler, Jonay Trénous, Marcos Treviso, Miguel Vera, and André F. T. Martins. 2019. OpenKiwi: An Open Source Framework for Quality Estimation. In *Proc. of the Annual Meeting of the Association for Computational Linguistics, System Demonstration*. Association for Computational Linguistics.
- Hyun Kim, Jong-Hyeok Lee, and Seung-Hoon Na. 2017. Predictor-Estimator using Multilevel Task Learning with Stack Propagation for Neural Quality Estimation. In *Conference on Machine Translation (WMT)*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M Rush. 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation. *arXiv preprint arXiv:1701.02810*.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual Language Model Pretraining. *arXiv preprint arXiv:1901.07291*.
- António V. Lopes, M. Amin Farajian, Gonçalo Correia, Jonay Trénous, and André F. T. Martins. 2019. Unbabel’s Submission to the WMT2019 APE Shared Task: Bert encoder-decoder for Automatic Post-Editing. In *Under review*.
- André F. T. Martins, Miguel B. Almeida, and Noah A. Smith. 2013. Turning on the Turbo: Fast Third-Order Non-Projective Turbo Parsers. In *Proc. of the Annual Meeting of the Association for Computational Linguistics*.
- André F. T. Martins, Ramon Astudillo, Chris Hokamp, and Fábio Kepler. 2016. Unbabel’s Participation in the WMT16 Word-Level Translation Quality Estimation Shared Task. In *Conference on Machine Translation (WMT)*.
- André F. T. Martins, Marcin Junczys-Dowmunt, Fabio Kepler, Ramon Astudillo, Chris Hokamp, and Roman Grundkiewicz. 2017. Pushing the Limits of Translation Quality Estimation. *Transactions of the Association for Computational Linguistics (to appear)*.
- Matteo Negri, Marco Turchi, Rajen Chatterjee, and Nicola Bertoldi. 2018. eSCAPE: a Large-scale Synthetic Corpus for Automatic Post-Editing. *arXiv preprint arXiv:1803.07274*.
- Franz Josef Och. 2003. Minimum Error Rate Training in Statistical Machine Translation. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 160–167. Association for Computational Linguistics.
- Michael J. D. Powell. 1964. An Efficient Method for Finding the Minimum of a Function of Several Variables without Calculating Derivatives. *The Computer Journal*, 7(2):155–162.
- Lucia Specia, Carolina Scarton, and Gustavo Henrique Paetzold. 2018. Quality Estimation for Machine Translation. *Synthesis Lectures on Human Language Technologies*, 11(1):1–162.
- Jiayi Wang, Kai Fan, Bo Li, Fengming Zhou, Boxing Chen, Yangbin Shi, and Luo Si. 2018. Alibaba Submission for WMT18 Quality Estimation Task. In *Conference on Machine Translation (WMT)*.