

# eSCAPE: a Large-scale Synthetic Corpus for Automatic Post-Editing

Matteo Negri<sup>(1)</sup>, Marco Turchi<sup>(1)</sup>, Rajen Chatterjee<sup>(1,2)</sup>, Nicola Bertoldi<sup>(1)</sup>

<sup>(1)</sup>Fondazione Bruno Kessler, <sup>(2)</sup>University of Trento

Trento, Italy

{negri, turchi, chatterjee, bertoldi}@fbk.eu

## Abstract

Training models for the automatic correction of machine-translated text usually relies on data consisting of (*source*, *MT*, *human-post-edit*) triplets providing, for each source sentence, examples of translation errors with the corresponding corrections made by a human post-editor. Ideally, a large amount of data of this kind should allow the model to learn reliable correction patterns and effectively apply them at test stage on unseen (*source*, *MT*) pairs. In practice, however, their limited availability calls for solutions that also integrate in the training process other sources of knowledge. Along this direction, state-of-the-art results have been recently achieved by systems that, in addition to a limited amount of available training data, exploit artificial corpora that approximate elements of the “gold” training instances with automatic translations. Following this idea, we present eSCAPE, the largest freely-available Synthetic Corpus for Automatic Post-Editing released so far. eSCAPE consists of millions of entries in which the *MT* element of the training triplets has been obtained by translating the *source* side of publicly-available parallel corpora, and using the *target* side as an artificial human post-edit. Translations are obtained both with phrase-based and neural models. For each MT paradigm, eSCAPE contains 7.2 million triplets for English–German and 3.3 millions for English–Italian, resulting in a total of 14,4 and 6,6 million instances respectively. The usefulness of eSCAPE is proved through experiments in a general-domain scenario, the most challenging one for automatic post-editing. For both language directions, the models trained on our artificial data always improve MT quality with statistically significant gains. The current version of eSCAPE can be freely downloaded from: <http://hltshare.fbk.eu/QT21/eSCAPE.html>.

**Keywords:** Automatic Post-editing, Machine Translation

## 1. Introduction

Automatic post-editing (APE) for machine translation (MT) aims to fix recurrent errors made by the MT decoder by learning from correction examples. As a post-processing step, APE has several possible applications, especially in *black-box* scenarios (*e.g.* when working with a third-party translation engine) in which the MT system is used “as is” and is not directly accessible for retraining or for more radical internal modifications. In such scenarios, as pointed out by Chatterjee et al. (2015), APE systems can help to: *i*) improve MT output by exploiting information unavailable to the decoder, or by performing a deeper text analysis that is too expensive at the decoding stage; *ii*) provide professional translators with improved MT output quality to reduce (human) post-editing effort, and *iii*) adapt the output of a general-purpose MT system to the lexicon/style requested in a specific application domain.

The training of APE systems usually relies on data sets comprising (*source*, *MT*, *human-post-edit*) triplets, in which the *source* sentence in a given language has been automatically translated to produce the *MT* element that, in turn, has been manually corrected to produce the *human-post-edit*. In this supervised learning setting, the goal is to learn from the training data (and possibly generalise) the appropriate corrections of systematic errors made by the MT system, and apply them at test stage on unseen (*source*, *MT*) pairs. BLEU (Papineni et al., 2002) and TER (Snover et al., 2006) computed against reference human post-edits are the standard evaluation metrics for the task. Their respective improvements and reductions are usually compared against the baseline scores obtained by the orig-

inal MT output that has been left untouched (*i.e.* raw, non post-edited translations).

Early works on this problem date back to (Allen and Hogan, 2000; Simard et al., 2007), which addressed the problem as a “monolingual translation” task in which raw MT output in the target language has to be translated, in the same language, into a fluent and adequate translation of the original source text. Although the general monolingual translation approach to the problem is still the same, over the years the proposed solutions evolved in several ways, first by refining the decoding approach and then, in the last couple of years, by radically changing the core APE technology.

Decoding refinements successfully explored, for instance, the integration of *source* information for enhanced (joint, context-aware) input representation, either in the standard phrase-based MT (PBMT) framework (Béchara et al., 2011) or in more elegant batch factored models (Chatterjee et al., 2016) and online PBMT models (Chatterjee et al., 2017b). More recently, radical paradigm changes followed the “neural revolution” witnessed in the MT field. The current state of the art is indeed represented by single/multi-source neural APE systems, the former relying on the log-linear combination of monolingual and bilingual models (Junczys-Dowmunt and Grundkiewicz, 2016), and the latter learning from source and target information in a joint fashion (Chatterjee et al., 2017a). Recent works addressed the problem by also integrating external information such as word-level quality estimation scores (Chatterjee et al., 2017c) as a way to guide neural APE decoding towards better corrections.

Unsurprisingly, the impressive gains achieved by the neu-

ral solutions come at the cost of a much higher data demand compared to the PBMT methods. To overcome this problem, the latest published results on neural APE have been obtained by exploiting synthetically-created data during training (Junczys-Dowmunt and Grundkiewicz, 2016; Junczys-Dowmunt and Grundkiewicz, 2017; Variš and Bojar, 2017; Hokamp, 2017; Chatterjee et al., 2017a).

These trends, which emerged after three rounds of the APE task organised within the Conference on Machine Translation (WMT) (Bojar et al., 2015; Bojar et al., 2016; Bojar et al., 2017), clearly indicate that: *i*) information from the *source* text is definitely useful to train reliable APE models, and *ii*) the limited availability of “gold” training corpora made of (*source*, *MT*, *human-post-edit*) triplets calls for workarounds to unleash the full potential of state-of-the-art but data-demanding neural systems.

The eSCAPE corpus presented in this paper meets such demand by providing APE research with a large-scale synthetically-created data set consisting of millions of triplets for two language pairs: English–German and English–Italian. Starting from a collection of publicly-available parallel corpora, it was built by automatically translating the *source* element of each sentence pair both with phrase-based and neural MT models, and using the original *MT* element as representative of a possible human correction.

This paper reports on the initial part of a roadmap aiming at more ambitious objectives. Future releases of the corpus will indeed include larger volumes of instances (translated with both MT paradigms) covering a larger spectrum of language combinations. The following sections provide an overview of the existing resources (Section 2), a description of eSCAPE (Section 3) and a discussion of experiments with the corpus (Section 4).

## 2. Related Work: Existing APE Corpora

The growing interest towards APE has to confront with the hard truth of data scarcity. Although nowadays post-edited data are a clear by-product of industrial translation workflows, the largest part of the daily work done by professional translators focuses on proprietary or copyright data that cannot be released. Though present in the industrial sector (as confirmed by recent works coming from big players like SYSTRAN (Crego et al., 2016) or eBay (Mathur et al., 2017)), APE technology is still more a matter of in-house development rather than a framework motivating free data sharing.

The few existing corpora that are usable for APE research can be classified into two types: *i*) the aforementioned “gold” data sets made of (*source*, *MT*, *human-post-edit*) triplets, and *ii*) the synthetic ones, to which our eSCAPE corpus belongs, in which some elements of the triplets derive from automatic translation. The remainder of this section provides an inventory of the existing APE corpora. As also shown in Table 1, the global picture is quite fragmented, with domain-specific data sets covering different language pairs, containing different types of post-edits and, most importantly, usually featuring a relatively small size.

### 2.1. “Gold” corpora

The Autodesk Post-Editing Data corpus (Zhechev, 2012)<sup>1</sup> is one of such resources. It mainly covers the domain of software user manuals, with English sentences translated with Autodesk’s in-house PBMT system into several languages (simplified and Traditional Chinese, Czech, French, German, Hungarian, Italian, Japanese, Korean, Polish, Brazilian Portuguese, Russian, Spanish) with between 30,000 and 410,000 segments per language. Post-edits are made by professional translators.

Part of the Autodesk corpus has been used by Chatterjee et al. (2015) to compare different APE techniques in a controlled setting. For six target languages (Czech, German, Spanish, French, Italian and Polish), this subset comprises around 16,000 (*source*, *MT*, *human-post-edit*) triplets that share the same English *source*. To ease the replicability of their experiments and the reuse of the selected triplets, the authors released the scripts used for data extraction.<sup>2</sup>

Another useful resource is described in (Potet et al., 2012).<sup>3</sup> It consists of 10,881 triplets in which a French source sentence taken from several news corpora is translated into English by a PBMT system. Post-edits were collected using Amazon Mechanical Turk following strict control reviewing procedures to guarantee correction quality.

Two smaller corpora are respectively described in (Specia et al., 2010) and (Specia, 2011). The former consists of 4,000 English sentences from Europarl (Koehn, 2005), which were translated into Spanish by a PBMT system and manually post-edited by professional translators. The latter, which covers the news domain, includes 2,525 French–English PBMT translations and 1,000 English–Spanish translations with professional post-edits.

Other useful data have been released by the organisers of the WMT APE task. The first round of the task (Bojar et al., 2015) presented participants with around 12,000 English–Spanish training data drawn from the news domain, with translations derived from a PBMT system. A peculiarity of this corpus is that post-edits were collected from a non-professional crowdsourced workforce, with possible drops in terms of reliability and consistency.<sup>4</sup>

The second round of the task (Bojar et al., 2016) presented participants with a corpus released within the EU project QT21,<sup>5</sup> the same used for the WMT’16 quality estimation task. It comprises 13,000 English–German training data drawn from the information technology (IT) domain, with source sentences translated by a PBMT system and post-edits collected from professional translators. The combination of domain specificity and higher post-editing quality resulted in significant gains over the baseline.

<sup>1</sup><https://autodesk.app.box.com/Autodesk-PostEditing>

<sup>2</sup><https://bitbucket.org/turchmo/apeatfbk/src/master/papers/ACL2015/>

<sup>3</sup><http://www-clips.imag.fr/geod/User/marion.potet/index.php>

<sup>4</sup>This is a possible cause of the poor results achieved by participants: none of them was indeed able to beat the APE task baseline represented by a “do-nothing” system that leaves all the raw MT translations unmodified.

<sup>5</sup><http://www.qt21.eu/>

Corpus	Type	Languages	Domain	Size	Post-edits
(Specia et al., 2010)	Gold	En-Es	LEGAL	4K	Professional
(Specia, 2011)	Gold	Fr-En/En-Es	NEWS	2.5K/1K	Professional
(Zhechev, 2012)	Gold	En-Ch/Cs/Ff/De/Hu/It/Ja/Ko/Pl/Br/Pt/Ru/Es	IT	30K-410K	Professional
(Potet et al., 2012)	Gold	Fr-En	NEWS	11K	Crowd
(Bojar et al., 2015)	Gold	En-Es	NEWS	12K	Crowd
(Bojar et al., 2016)	Gold	En-De	IT	13K	Professional
(Bojar et al., 2017)	Gold	En-De/De-En	IT/MEDICAL	13K/26K	Professional
(Junczys-Dowmunt and Grundkiewicz, 2016)	Artificial	En-De	IT	4.3M	-

Table 1: Inventory of existing APE corpora

The third round of the task (Bojar et al., 2017) focused on both English–German and German–English data (also in this case provided by the QT21 project (Specia et al., 2017) and shared with the WMT’17 quality estimation task). English–German training data are drawn from the IT domain and consist of around 13,000 triplets. German–English training data, instead, come from the medical domain and comprise around 26,000 triplets. In both cases, translations were produced by a customised PBMT system and post-edited by professional translators.

## 2.2. Synthetic corpora

The use of synthetic resources aims to overcome the aforementioned problem of “gold” data scarcity with approximate solutions. This can be done in different ways. Several previous works have shown the viability of mimicking the ideal scenario in which the training triplets include actual human post-edits of machine-translated text by learning, instead, from the weaker connection between the MT output and external references. Though with variable margins, (Oflazer and El-Kahlout, 2007; Béchara et al., 2011; Rubino et al., 2012) report translation quality improvements in the PBMT scenario with post-editing components trained on (*source*, *MT*, *reference*) triplets. To the best of our knowledge, though potentially useful to APE research, none of such previous works released reusable datasets. When moving to the data-demanding neural framework, data scarcity becomes a major problem that definitely calls for the external support of artificial corpora that are orders of magnitude larger than the current training sets.

A widely used resource, described in (Junczys-Dowmunt and Grundkiewicz, 2016), was included in the training set of the winning (and almost all) submissions to the last two English–German rounds of the APE task at WMT (IT domain). It consists of 4.3 million instances created by first filtering a subset of IT-related sentences from the German Common Crawl corpus<sup>6</sup>, and then by using two English–German and German–English PBMT systems trained on in-domain IT corpora for a round-trip translation of the selected sentences (De→En→De). The final triplets were created by using: *i*) the English translations as (artificial) *source* sentences, *ii*) the round-trip German translations as (artificial) uncorrected *MT* output, and *iii*) the original German sentences as (artificial) *post-edits*.

By construction, this artificial data set approximates the quality of gold corpora by trying to keep a weak connection between the “post-edits” and the MT output. Keeping such connection, however, comes at the cost of having two levels of potential noise in the data, namely the possible errors introduced by the German–English translation needed to produce the *source* element of each triplet, and those of the English–German translations performed to produce the *MT* output. The approach we adopted to create the eSCAPE corpus, instead, follows a different strategy. As described in the next section, we start from parallel data and perform one single automatic translation step to produce the *MT* element of our triplets. The connection between MT output and “post-edits” is hence weaker than in (Junczys-Dowmunt and Grundkiewicz, 2016) due to the fact that our “post-edits” are actually independent reference translations of the source sentences. However, the possible noise introduced by translation errors can only affect one element of our triplets. Analysing the trade-off between translation noise and *MT-post-edits* proximity is out of the scope of this work but it is definitely an interesting aspect for future investigations.

## 3. The eSCAPE Corpus

The eSCAPE corpus<sup>7</sup> consists in two datasets (En-De and En-It) made of (*source*, *MT*, *reference*) triplets, where the *MT* segment is obtained by translating the *source* both with phrase-based and neural MT models. Its creation started from parallel (*source*, *target*) data collected from the WEB by merging several corpora belonging to various domains. Table 2 lists all the corpora used, indicating their domain and size in terms of number of sentences. Since some data sets, such as PatTR and Common Crawl, are only available in one language pair (En-De), the total number of sentences is different between the two language directions (En-De is twice larger than En-It). Apart from PatTR<sup>8</sup> and Common Crawl, all the datasets are available in the OPUS repository<sup>9</sup>. Before building the translation systems and producing the *MT* segments, all the corpora reported in Table 2 have been concatenated and shuffled (to avoid blocks of sentences belonging to the same domain) removing duplicates. This resulted in 7,258,533

<sup>7</sup><http://hltshare.fbk.eu/QT21/eSCAPE.html>

<sup>8</sup><http://www.cl.uni-heidelberg.de/statnlpgroup/pattr/>

<sup>9</sup><http://opus.lingfil.uu.se>

<sup>6</sup>[commoncrawl.org](http://commoncrawl.org)

English–German and 3,357,371 English–Italian sentence pairs.

Corpus	Domain	En-De	En-It
Europarl v7	LEGAL	1,920,209	1,909,115
ECB	LEGAL	11,317	193,047
Common Crawl	MIXED	2,399,123	-
JRC Acquis	LEGAL	719,372	810,979
News Commentary v11	NEWS	242,770	40,009
Ted Talks	MIXED	143,836	181,874
EMEA	MEDIC.	1,108,752	1,081,134
PatTR in-domain	MEDIC.	1,848,303	-
Wikipedia Titles	MEDIC.	10,406	-
Gnome	IT	28,439	319,141
Ubuntu	IT	13,245	21,014
KDE4	IT	224,035	175,058
OpenOffice	IT	42,391	-
PHP	IT	39,707	35,538
TOTAL		8,853,762	4,128,128

Table 2: List of datasets merged in the eSCAPE corpus.

### 3.1. MT systems

Driven by the need of translating these large quantities of source segments, the ModernMT toolkit (Bertoldi et al., 2017) has been used as translation system to generate both phrase-based and neural outputs. ModernMT is a new open-source MT software that consolidates the current state-of-the-art MT technology into a single and easy-to-use product. The toolkit adapts to the context in real-time and is capable of learning from (and evolving through) interaction with users, with the final aim of increasing MT-output utility for the translator in a real professional environment.

To avoid the risk of translating source segments that are in the training set, the collected sentence pairs were split in 4 slices: 3 parts were used to train the MMT models and the remaining one was translated. In this cross-validation setting, one sentence pair has been processed once for each experiment, either in training or in test.

For **phrase-based MT**, ModernMT uses high-performance embedded databases to store parallel and monolingual language data and associated statistics. Instead of pre-computing phrase-based feature function scores, these are computed on the fly, at translation time, from raw statistics. This allows the MMT toolkit to significantly speed up the training and test processes, to easily scale to large quantities of data, and to adapt on-the-fly to new domains. Training and test of the phrase-based models were run in parallel on several CPUs for around one week. Final performance, computed on a subset of the data, is 36.76 BLEU points for English–German and 38.08 for English–Italian.

For **neural MT**, the toolkit builds on the extension of a generic neural MT system based on the Nematus toolkit (Sennrich et al., 2017)<sup>10</sup>, implementing the encoder-decoder-attention model architecture by (Bahdanau et al., 2014). Such extension consists in an internal dynamic

memory, storing external user translation memories (TMs). When ModernMT receives a translation query, it quickly analyses its context, recalls from its memory the most related translation examples, and instantly adapts its neural network to the query (Farajian et al., 2017). Training and test of the neural models were run on one GPU (NVIDIA Tesla K80) for around three weeks. Final performance is 38.17 BLEU points for English–German and 41.01 for English–Italian.

To give the possibility for experiments on domain-adaptation for APE, each eSCAPE triplet is associated to a label indicating the name of the corpus from which the original (*source*, *reference*) pair was extracted.

## 4. Experiments

To test the usefulness of the eSCAPE corpus, we run APE experiments for both the language pairs covered by the data set. En-De and En-It data were first tokenised and then split into dev (2,000 triplets), test (10,000) and training (the remaining instances). For the sake of comparison, we performed the same data splits for both the phrase-based and for the neural-based section of the corpus.

As APE system, we chose the best system at this year’s round of the WMT APE shared task (Chatterjee et al., 2017a). It consists in a neural multi-source model, in which the source and the MT segment are encoded separately and then merged together by a feed-forward network layer. A shared dropout is applied to both source and MT encoders. In this multi-source architecture both the encoders are trained jointly.

In our experiments, the hyper-parameters of all the systems in both language directions were the same. The vocabulary was created by selecting 50,000 most frequent sub-words, following the BPE approach of Sennrich et al. (2016b). Word-embedding and GRU (gated recurrent unit) hidden-state sizes were both set to 1024. Network parameters were optimized with Adagrad (Duchi et al., 2011) with a learning rate of 0.01. Source and target dropout was set to 10%, whereas encoder and decoder hidden states, weighted source context, and embedding dropout was set to 20% (Sennrich et al., 2016a). The batch size was set to 100 samples, with a maximum sentence length of 50 sub-words.

During training, the performance of the APE system was monitored on the development set and, at the end of the training phase, the model with highest BLEU score was used to post-edit the test set. The results are reported in Table 3 where, for both the language pairs and for both phrase-based and neural-based artificial data, the performance of the APE systems is compared against the “*do-nothing*” APE baseline (i.e. a system that leaves all the raw MT output unmodified).

It is interesting to note that APE systems outperform the baselines in both language settings with statistically significant gains.<sup>11</sup> This holds true both when they are trained

<sup>11</sup>Although we consider the measured gains as a good indicator of the usefulness of using eSCAPE for training APE models, a study involving human evaluation would allow us to draw definite conclusions. Such a costly study, however, falls out of the scope of this paper and is left for future work.

<sup>10</sup><https://github.com/rsennrich/nematus>

	En-De	En-It
Phrase-based MT		
Do-Nothing baseline	36.76	38.08
APE	<b>38.15</b>	<b>39.80</b>
Neural MT		
Do-Nothing baseline	38.17	41.01
APE	<b>39.21</b>	<b>42.15</b>

Table 3: Neural APE results (BLEU score improvements are statistically significant with  $p < 0.05$  computed with paired bootstrap resampling (Koehn, 2004)).

and tested on artificial data built from phrase-based models (+1.39 on En-De, +1.72 on En-It), and when training and test are performed on artificial data derived from neural models (+1.04 on En-De, +1.14 on En-It).<sup>12</sup>

The observed gains vary for the two language pairs (with highest results on En-It) and depending on the type of data used. Concerning this latter aspect, the higher quality of neural MT output results in lower gains on both language settings. This confirms previous outcomes from the WMT APE task: the higher the baseline (i.e. the BLEU score of the raw MT output), the lower the number of correction patterns that can be learned from the training data and the possibility to leverage their applicability to test data (Bojar et al., 2017).

Differently from the most recent shared evaluation settings (i.e. WMT’16 and WMT’17), in which neural APE has been tested in narrow domains, our results indicate that APE systems trained on the eSCAPE corpus can be also effective in the more challenging mixed-domain condition, where the correction rules are sparse across different domains, hence difficult to be learned and generalized. Considering the negative outcomes of the WMT’15 pilot task, which proposed a challenging evaluation setting based on general news data in which none of the participants was able to beat the “do-nothing” baseline, this is an interesting finding that adds value to our resource.

The BLEU score improvements also confirm the findings of (Oflazer and El-Kahlout, 2007; Béchara et al., 2011; Rubino et al., 2012) and extend them to neural APE. In fact, they report translation quality improvements in the PBMT scenario with an APE trained on *source*, *MT*, and independent *reference*. This suggests that, despite the aforementioned weak connection between the MT output and the “post-edits” of our triplets, APE models can be effectively trained on the eSCAPE corpus.

## 5. Conclusion and Future Work

We presented the eSCAPE corpus, a large-scale Synthetic Corpus for Automatic Post-Editing consisting of millions of (*source*, *MT*, *post-edit*) triplets created via machine translation. eSCAPE is designed to support the recent trends in automatic post-editing, which show a clear predominance of data-demanding neural approaches. To cope

<sup>12</sup>Though interesting, other settings in which the two sections of eSCAPE are either combined or alternatively used one for training and one for test fall out of the scope of this paper and are left for future investigation.

with such demand, the current version of the corpus contains millions of triplets for two language pairs: English-German (14.4 millions) and English-Italian (6.6 millions). For both language pairs, half of the artificial data is obtained via phrase-based translation, while the other half is produced by better performing neural MT models. Having the same source sentences translated with both paradigms aims to enable future comparisons in the application of APE technology on the two types of output. The size of the corpus (the largest of its kind) is expected to ease model training and yield further state of the art improvements. Our preliminary experiments on mixed-domain data confirm this expectation: though trained on artificially-created instances, APE models significantly outperform baseline results in both language directions, independently from the MT technology underlying the data generation process. The work reported in this paper is the initial step of a more ambitious roadmap aimed to extend the resource with more data covering a larger spectrum of domains and language combinations. The current version of eSCAPE can be freely downloaded from: <http://hltshare.fbk.eu/QT21/eSCAPE.html>.

## Acknowledgments

This work has been partially supported by the EC-funded H2020 projects QT21 (grant agreement no. 645452) and ModernMT (grant agreement no. 645487). The authors gratefully acknowledge the support of NVIDIA Corporation with the donation of the Tesla K80 GPU used for this research.

## 6. Bibliographical References

- Allen, J. and Hogan, C. (2000). Toward the Development of a Post Editing Module for Raw Machine Translation Output: A Controlled Language Perspective. In *Proceedings of the Third International Controlled Language Applications Workshop (CLAW-2000)*, pages 62–71, Seattle, Washington, April.
- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Béchara, H., Ma, Y., and van Genabith, J. (2011). Statistical Post-Editing for a Statistical MT System. In *Proceedings of the 13th Machine Translation Summit*, pages 308–315, Xiamen, China, September.
- Bertoldi, N., Cattoni, R., Cettolo, M., Farajian, M., Federico, M., Caroselli, D., Mastrostefano, L., Rossi, A., Trombetti, M., Germann, U., et al. (2017). MMT: New Open Source MT for the Translation Industry. In *Proceedings of The 20th Annual Conference of the European Association for Machine Translation (EAMT)*, pages 86–91.
- Bojar, O., Chatterjee, R., Federmann, C., Haddow, B., Huck, M., Hokamp, C., Koehn, P., Logacheva, V., Monz, C., Negri, M., Post, M., Scarton, C., Specia, L., and Turchi, M. (2015). Findings of the 2015 Workshop on Statistical Machine Translation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 1–46, Lisbon, Portugal, September.

- Bojar, O., Chatterjee, R., Federmann, C., Graham, Y., Haddow, B., Huck, M., Jimeno Yepes, A., Koehn, P., Logacheva, V., Monz, C., Negri, M., Neveol, A., Neves, M., Popel, M., Post, M., Rubino, R., Scarton, C., Specia, L., Turchi, M., Verspoor, K., and Zampieri, M. (2016). Findings of the 2016 Conference on Machine Translation. In *Proceedings of the First Conference on Machine Translation*, pages 131–198, Berlin, Germany, August.
- Bojar, O., Chatterjee, R., Federmann, C., Graham, Y., Haddow, B., Huang, S., Huck, M., Koehn, P., Liu, Q., Logacheva, V., Monz, C., Negri, M., Post, M., Rubino, R., Specia, L., and Turchi, M. (2017). Findings of the 2017 Conference on Machine Translation (WMT17). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214, Copenhagen, Denmark, September.
- Chatterjee, R., Weller, M., Negri, M., and Turchi, M. (2015). Exploring the Planet of the APEs: a Comparative Study of State-of-the-art Methods for MT Automatic Post-Editing. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*, pages 156–161, July.
- Chatterjee, R., C. de Souza, J. G., Negri, M., and Turchi, M. (2016). The FBK Participation in the WMT 2016 Automatic Post-editing Shared Task. In *Proceedings of the First Conference on Machine Translation*, pages 745–750, Berlin, Germany, August. Association for Computational Linguistics.
- Chatterjee, R., Farajian, M. A., Negri, M., Turchi, M., Srivastava, A., and Pal, S. (2017a). Multi-source Neural Automatic Post-Editing: FBK’s participation in the WMT 2017 APE shared task. In *Proceedings of the Second Conference on Machine Translation*, pages 630–638, Copenhagen, Denmark, September.
- Chatterjee, R., Gebremelak, G., Negri, M., and Turchi, M. (2017b). Online Automatic Post-editing for MT in a Multi-Domain Translation Environment. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 525–535, Valencia, Spain, April.
- Chatterjee, R., Negri, M., Turchi, M., Federico, M., Specia, L., and Blain, F. (2017c). Guiding Neural Machine Translation Decoding with External Knowledge. In *Proceedings of the Second Conference on Machine Translation*, pages 157–168, Copenhagen, Denmark, September.
- Crego, J. M., Kim, J., Klein, G., Rebollo, A., Yang, K., Senellart, J., Akhanov, E., Brunelle, P., Coquard, A., Deng, Y., Enoue, S., Geiss, C., Johanson, J., Khalsa, A., Khiari, R., Ko, B., Kobus, C., Lorieux, J., Martins, L., Nguyen, D., Priori, A., Riccardi, T., Segal, N., Servan, C., Tiquet, C., Wang, B., Yang, J., Zhang, D., Zhou, J., and Zoldan, P. (2016). SYSTRAN’s Pure Neural Machine Translation Systems. *arXiv preprint 1610.05540*.
- Duchi, J., Hazan, E., and Singer, Y. (2011). Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, pages 2121–2159.
- Farajian, M. A., Turchi, M., Negri, M., Bertoldi, N., and Federico, M. (2017). Neural vs. Phrase-Based Machine Translation in a Multi-Domain Scenario. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 280–284, Valencia, Spain, April.
- Hokamp, C. (2017). Ensembling Factored Neural Machine Translation Models for Automatic Post-Editing and Quality Estimation. In *Proceedings of the Second Conference on Machine Translation*, pages 647–654, Copenhagen, Denmark, September.
- Junczys-Dowmunt, M. and Grundkiewicz, R. (2016). Log-linear Combinations of Monolingual and Bilingual Neural Machine Translation Models for Automatic Post-Editing. In *Proceedings of the First Conference on Machine Translation*, pages 751–758, Berlin, Germany, August. Association for Computational Linguistics.
- Junczys-Dowmunt, M. and Grundkiewicz, R. (2017). An Exploration of Neural Sequence-to-Sequence Architectures for Automatic Post-Editing. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, Taipei, Taiwan, November.
- Koehn, P. (2004). Statistical Significance Tests for Machine Translation Evaluation. In *Proceedings of the Empirical Methods on Natural Language Processing (EMNLP 2004)*, pages 388–395, Barcelona, Spain.
- Koehn, P. (2005). Europarl: A Parallel Corpus for Statistical Machine Translation. In *Proceedings of the Tenth Machine Translation Summit (MT Summit X)*, pages 79–86, Phuket, Thailand, September.
- Mathur, P., Ueffing, N., and Leusch, G. (2017). Generating Titles for Millions of Browse Pages on an e-Commerce Site. In *Proceedings of the 10th International Conference on Natural Language Generation (INLG2017)*, pages 158–167, Santiago de Compostela, Spain, September.
- Oflazer, K. and El-Kahlout, I. D. (2007). Exploring Different Representational Units in English-to-Turkish Statistical Machine Translation. In *Proceedings of the Second Workshop on Statistical Machine Translation*, StatMT ’07, pages 25–32, Prague, Czech Republic.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL ’02, pages 311–318, Philadelphia, Pennsylvania.
- Potet, M., Esperança-Rodier, E., Besacier, L., and Blanchon, H. (2012). Collection of a Large Database of French-English SMT Output Corrections. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012*, pages 4043–4048, Istanbul, Turkey.
- Rubino, R., Huet, S., Lefèvre, F., and Linares, G. (2012). Statistical Post-Editing of Machine Translation for Domain Adaptation. In *Proceedings of the European Association for Machine Translation (EAMT)*, pages 221–228, Trento, Italy, May.
- Sennrich, R., Haddow, B., and Birch, A. (2016a). Edinburgh Neural Machine Translation Systems for WMT 16. In *Proceedings of the First Conference on Machine Translation*, pages 371–376, Berlin, Germany, August.
- Sennrich, R., Haddow, B., and Birch, A. (2016b). Neural Machine Translation of Rare Words with Subword Units.

- In *Proceedings of the 54th Annual Meeting on Association for Computational Linguistics*, pages 1715–1725, Berlin, Germany.
- Sennrich, R., Firat, O., Cho, K., Birch, A., Haddow, B., Hirschler, J., Junczys-Dowmunt, M., Läubli, S., Miceli Barone, A. V., Mokry, J., and Nadejde, M. (2017). Nematius: a Toolkit for Neural Machine Translation. In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain, April.
- Simard, M., Goutte, C., and Isabelle, P. (2007). Statistical Phrase-Based Post-Editing. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 508–515, Rochester, New York, April.
- Snover, M., Dorr, B., Schwartz, R., Micciulla, L., and Makhoul, J. (2006). A Study of Translation Edit Rate with Targeted Human Annotation. In *Proceedings of the Association for Machine Translation in the Americas*, pages 223–231, Cambridge, Massachusetts, August.
- Specia, L., Cancedda, N., and Dymetman, M. (2010). A Dataset for Assessing Machine Translation Evaluation Metrics. In *Seventh Conference on International Language Resources and Evaluation, LREC 2010*, pages 3375–3378, Valletta, Malta.
- Specia, L., Harris, K., Blain, F., Burchardt, A., Mackentanz, V., Skadina, I., Negri, M., and Turchi, M. (2017). Translation Quality and Productivity: A Study on Rich Morphology Languages. In *Proceedings of the 16th Machine Translation Summit*, pages 55–71, Nagoya, Japan, September.
- Specia, L. (2011). Exploiting Objective Annotations for Measuring Translation Post-editing Effort. In *15th Conference of the European Association for Machine Translation, EAMT*, pages 73–80, Leuven, Belgium.
- Variš, D. and Bojar, O. (2017). CUNI System for WMT17 Automatic Post-Editing Task. In *Proceedings of the Second Conference on Machine Translation*, pages 661–666, Copenhagen, Denmark, September.
- Zhechev, V. (2012). Machine Translation Infrastructure and Post-Editing Performance at Autodesk. In *AMTA 2012 Workshop on Post-Editing Technology and Practice (WPTP 2012)*, pages 87–96, San Diego, CA, USA.