

Incremental Adaptation of NMT for Professional Post-editors: A User Study

Miguel Domingo¹ and Mercedes García-Martínez² and Álvaro Peris¹ and Alexandre Helle² and Amando Estela² and Laurent Bié² and Francisco Casacuberta¹ and Manuel Herranz²

¹PRHLT Research Center - Universitat Politècnica de València
{midobal, lvapeab, fcn}@prhl.upv.es

²Pangeanic / B.I Europa - PangeaMT Technologies Division
{m.garcia, a.helle, a.estela, l.bie, m.herranz}@pangeanic.com

Abstract

A common use of machine translation in the industry is providing initial translation hypotheses, which are later supervised and post-edited by a human expert. During this revision process, new bilingual data are continuously generated. Machine translation systems can benefit from these new data, incrementally updating the underlying models under an online learning paradigm. We conducted a user study on this scenario, for a neural machine translation system. The experimentation was carried out by professional translators, with a vast experience in machine translation post-editing. The results showed a reduction in the required amount of human effort needed when post-editing the outputs of the system, improvements in the translation quality and a positive perception of the adaptive system by the users.

1 Introduction

Translation post-editing is a common use case of machine translation (MT) in the industrial environment. Post-editing consists of the supervision by a human agent of outputs generated by an MT system, who corrects the errors made by the MT system. As MT systems are continuously improving their capabilities, translation post-editing has acquired major relevance in the translation market (Arenas, 2008; Hu and Cadwell, 2016). As a byproduct of this process, new data are continuously generated. These data have valuable properties: they are domain-specific training samples,

© 2019 The authors. This article is licensed under a Creative Commons 4.0 licence, no derivative works, attribution, CC-BY-ND.

which can be leveraged for adapting the system towards a given domain or post-editor. Moreover, an adaptive system can learn from its mistakes. In other words, it can avoid making the same errors again.

A typical way of profiting from these post-edits consists in updating the system following an online learning paradigm: as the user validates a post-edit, the system is incrementally updated, by taking into account this sample. Hence, when the system generates the next translation, it will consider the previous user post-edits. It is expected that better translations (or more suited to the human post-editor preferences) will be produced.

In this paper, we evaluate this strategy in an industrial scenario. We study the enhancements brought about by an adaptive system via online learning, and the effects on the post-editing process of data generated by a neural machine translation (NMT) system. To that end, we firstly evaluate our system under laboratory conditions. Next, we conduct the evaluation of the system on a production environment. This experiment involved professional translators, who regularly rely on MT post-editing in their workflow. The results show improvements of adaptive systems in terms of productivity and translation quality.

2 Related work

Translation post-editing has been a widely adopted practice in the industry for a long time (e.g., Vasconcellos and León, 1985). As MT technology advanced and improved, the post-editing process gained more relevance and many user studies have demonstrated its capabilities (Aziz et al., 2012; Bentivogli et al., 2016; Castilho et al., 2017; Green et al., 2013a).

Adapting an MT system from user post-edits via

online learning techniques has also attracted the attention of researchers and industry parallel to the rise of the post-editing protocol. Many advances in this direction were achieved during the CasMaCat (Alabau et al., 2013) and MateCat (Federico et al., 2014) projects, which adapted phrase-based statistical machine translation systems incrementally from user post-edits.

Following recent breakthroughs in NMT technology, some works studied the construction of adaptive systems via online learning in this post-editing scenario. Turchi et al. (2017) and Peris et al. (2017) proposed to adapt an NMT system with post-edited samples to a new domain via online learning. Other works aimed to refine these adaptation techniques: Wuebker et al. (2018) applied sparse updates; Kothur et al. (2018) introduced a dictionary of translations for dealing with the novel words included in the new domain. However, in all these works, the users were simulated, due to the economical costs of involving humans within experiments.

User studies on online adaptation from post-edits have been conducted, mainly for phrase-based statistical machine translation systems (Alabau et al., 2016; Bentivogli et al., 2016; Denkowski et al., 2014; Green et al., 2013b). Regarding the NMT technology, several user studies have been recently conducted, analyzing different MT technologies (Koponen et al., 2019; Jia et al., 2019) or protocols (Daems and Macken, 2019). The closest work to ours was developed by Karimova et al. (2018), who showed savings in human effort, due to the effect of online learning. But in contrast to our work, the individuals used in Karimova et al. (2018) were students, whereas we conducted the study using professional, experienced translators.

3 Online learning from NMT post-edits

NMT relies on the statistical formalization of MT (Brown et al., 1990). The goal is to obtain, given a source sentence \mathbf{x} , its most likely translation $\hat{\mathbf{y}}$:

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} \Pr(\mathbf{y} | \mathbf{x}) \quad (1)$$

This probability is directly modeled by a neural network with parameters Θ :

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} \log p(\mathbf{y} | \mathbf{x}; \Theta) \quad (2)$$

This neural network usually follows an encoder–decoder architecture, featuring recurrent (Bahdanau et al., 2015; Sutskever et al., 2014) or convolutional networks (Gehring et al., 2017) or attention mechanisms (Vaswani et al., 2017). The parameters of the model are typically estimated jointly on large parallel corpora, via stochastic gradient descent (SGD; Robbins and Monro, 1951; Rumelhart et al., 1986). At decoding time, the system obtains the most likely translation by means of a beam search method.

3.1 Adaption from post-edits via online learning

During the usage of the MT system, we can leverage the post-edited samples for continuously adapting the system on the fly, as soon as a sentence has been post-edited. This procedure is described in Algorithm 1: for each sentence to be translated (\mathbf{x}), the system produces a translation hypothesis $\hat{\mathbf{y}}$. The user post-edits this sentence, obtaining a corrected version of it (\mathbf{y}). Right after this post-editing process, and before translating the next sample, the NMT system is updated, taking into account \mathbf{x} and \mathbf{y} .

```

Input :  $\Theta_1$  (initial NMT system),
          $\{\mathbf{x}_n\}_{n=1}^N$  (source sentences)
1 begin
2    $n \leftarrow 1$ 
3   while  $n \leq N$  do
4      $\hat{\mathbf{y}}_n \leftarrow \text{Translate}(\mathbf{x}_n, \Theta_n)$ 
5      $\mathbf{y}_n \leftarrow \text{Post-edit}(\mathbf{x}_n, \hat{\mathbf{y}}_n)$ 
6      $\Theta_{n+1} \leftarrow \text{Update}((\mathbf{x}_n, \mathbf{y}_n), \Theta_n)$ 
7      $n \leftarrow n + 1$ 

```

Algorithm 1: Adaptation via online learning during NMT post-editing.

This adaptation of the NMT model can be performed following the same method used in regular training: SGD.

4 Experimental framework

We now describe the experimental conditions arranged in our study: the translation systems and environment, the main features of the tasks under study and the evaluation criteria considered.

4.1 NMT systems

Our NMT system was a recurrent encoder–decoder with an additive attention mechanism (Bahdanau et al., 2015), built with OpenNMT-py (Klein et al., 2017). We used long short-term memory units

(Gers et al., 2000) and we set all model dimensions to 512. The system was trained using Adam (Kingma and Ba, 2014) with a fixed learning rate of 0.0002 (Wu et al., 2016) and a batch size of 60. We applied label smoothing of 0.1 (Szegedy et al., 2015). At the inference time, we used a beam search with a beam size of 6. We applied joint byte pair encoding to all corpora (Sennrich et al., 2016), using 32,000 merge operations.

The adaptive systems were built considering the findings from Peris and Casacuberta (2019), and conducting an evaluation on a development set. For each new post-edited sample, we performed two plain SGD updates, with a fixed learning rate of 0.05.

4.2 Translation environment

In order to assess the benefits of the adaptive system, we started by conducting an experiment with simulated users in a laboratory setting. This study is frequently carried out within the literature (e.g., Ortiz-Martínez, 2016), due to the economical costs of involving humans within experiments. Following common practices, we used the reference sentences as translation post-edits. Therefore, in the static scenario, we assessed the quality of the system using the references. In the adaptive scenario, we translated each source sentence and applied on-line learning with the corresponding reference.

Once we studied the behavior of the system under simulated conditions, we conducted the experiment with the real users. They were three professional translators, with an average of four years of experience, who regularly make use of MT in their workflow.

The experiment was conducted using SDL Trados Studio as the translation environment. This software is widely used in the translation industry, and all the participants use it in their daily work. Fig. 1 shows a screenshot of the SDL Trados Studio interface.

Our NMT system was deployed as a server, which delivered the translations to SDL Trados Studio and performed the adaptation using the post-edits. This system is compatible with all OpenNMT-py models and it is publicly available¹. We also developed a plugin that connected SDL Trados Studio with our systems.

¹<https://github.com/midobal/OpenNMT-py/tree/OnlineLearning>

4.3 Tasks and evaluation

We evaluated our systems on a real task from our production scenario. This task consisted in a small corpus belonging to a medico-technical domain (description of medical equipments), and was conformed by two documents of 150 sentences each, containing 1.7 and 2.7 thousand words respectively. The translation direction was from English to Spanish. Since we lacked an in-domain corpus, we trained a general system with the data from the translation task from WMT'13 (Bojar et al., 2013), consisting in 15 million parallel segments. Next, we applied the FDA data selection technique (Biçici and Yuret, 2015) for selecting related instances from our general corpus and a medical (UFAL, Bojar et al., 2017) and technological² ones. We selected 8 million additional segments, which were used for fine-tuning the general system.

The effects of adaptivity were assessed according to the post-editing time and to two common MT metrics: (h)BLEU (Papineni et al., 2002) and (h)TER (Snover et al., 2006). For ensuring consistent BLEU scores, we used sacreBLEU (Post, 2018). Since we computed per-sentence BLEU scores, we used exponential BLEU smoothing (Chen and Cherry, 2014). In order to determine whether two systems presented statistically significant differences, we applied approximate randomization tests (Riezler and Maxwell, 2005), with 10,000 repetitions and a p -value of 0.05.

5 Results

As introduced in the previous section, we first analyzed the adaptation process in a simulated environment. Next, we studied and discussed the results obtained in the user trials.

5.1 Adaptation with simulated users

Table 1 shows the results in terms of translation quality of a static system, compared with an adaptive one, updated using the reference samples. The results obtained on this synthetic setup support the usefulness of the adaptation via online learning: in all cases, the adaptive system achieved better TER and BLEU than the static one. These differences were statistically significant in all cases but one. We observed important gains in terms of TER (5.5

²<https://metashare.metanet4u.eu/go2/qtleapcorpus>

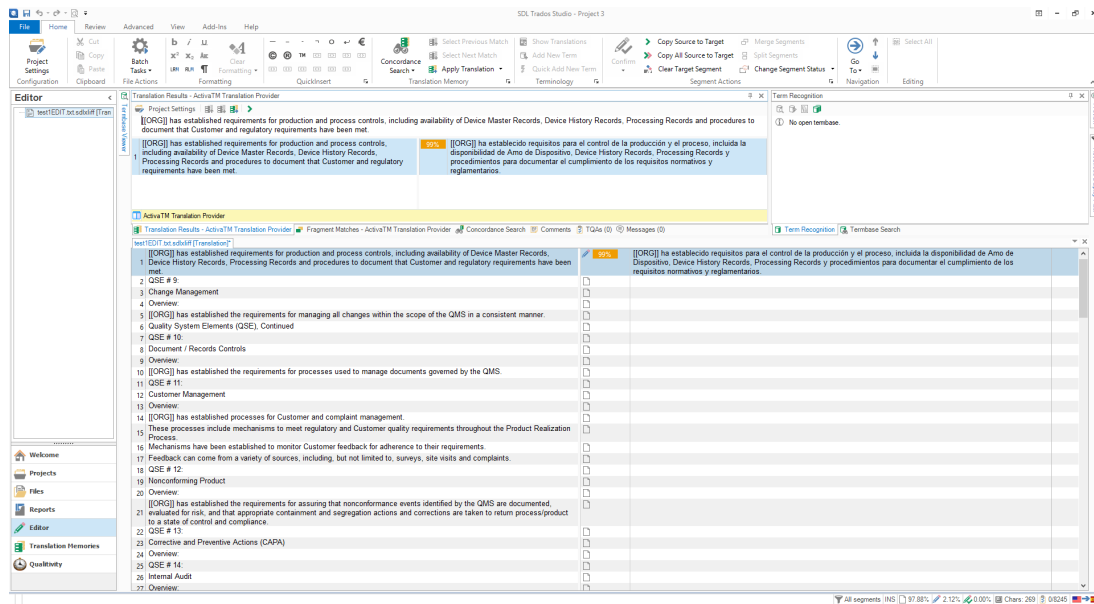


Figure 1: User Interface from SDL Trados Studio. From top to bottom, the first row and the leftmost column correspond to the user menus. On the next row, the middle column contains information about the segment that is being translated: on the left, the source sentence and, on the right, the MT translation. The right column displays the content of the terminological dictionary (if any). The document that is being translated appears on the bottom row: on the left, the original document and, on the right, the user post-edits.

and 1.1 points), which suggests a lower human effort required to for post-edit these samples. We also experimented with a larger document (1,500 sentences), belonging to the same domain. The adaptation to this larger document was more effective: we observed gains of 10.4 TER points and 13.6 BLEU points.

Test	System	TER [↓]	BLEU [↑]
T1	Static	54.0	26.9
	Adaptive	48.5 [†]	32.0 [†]
T2	Static	56.1	23.4
	Adaptive	55.0	26.3 [†]

Table 1: Results of the simulated experiments. Static systems stand for conventional post-editing, without adaptation. Adaptive systems refer to post-editing in an environment with online learning. TER and BLEU were computed against the reference sentences. [†] indicates statistically significant differences between the static and the adaptive systems.

Additionally to the assessment of the system in terms of translation quality, we need to satisfy an adequate latency, including decoding and updating times. Our NMT system was deployed in a CPU server, equipped with an Intel(R) Xeon(R) CPU E5-2686 v4 at 2.30GHz and 16GB of RAM. On average, generating a translation took the system 0.23 seconds and each update took 0.45 seconds. These low latencies allow a correct usage of the

system, as the flow of thoughts of the user remains uninterrupted (Nielsen, 1993).

5.2 Adaptation with human post-editors

User	Static	Adaptive
User 1	T1	T2
User 2	T2	T1
User 3	T1	T2

Table 2: Distribution of users (1, 2 and 3), test sets (T1 and T2) and scenarios (Static and Adaptive).

Once we tested our system in a simulated environment, we moved on to the experimentation with human post-editors. Three professional translators were involved in the experiment. For the adaptive test, all post-editors started the task with the same system, which was adapted to each user using their own post-edits. Therefore, at the end of the online learning process, each post-editor obtained a tailored system. For the static experiment, the initial NMT system remained fixed along the complete process. In order to avoid the influence of translating the same text multiple times, each participant post-edited a different test set under each scenario (static and adaptive), as shown in Table 2.

The main results of this experiment are shown in Table 3. These numbers are averages over the results obtained by the different post-editors. The




Figure 2: hBLEU per sentence of static and adaptive systems for both test sets (T1 and T2). Individual sentence scores are plotted for each system, static (red crosses) and adaptive (blue dots). The sentences were processed sequentially, hence, we can observe the progress of the system with its usage. To this end, we show a fit of the scores of each system, in dashed red and solid blue lines, for static and adaptive systems, respectively.

large reduction of post-editing time per sentence for the set T1 is especially relevant (an average of 7.5 seconds per sentence). In the test set T2, the post-editing time of the adaptive system was also slightly lower than the static system one, but only by 0.7 seconds.

Test	System	Time (s)	hTER [\downarrow]	hBLEU [\uparrow]
T1	Static	37.9	39.5	47.3
	Adaptive	30.4	34.2	55.1 [†]
T2	Static	45.8	38.4	45.7
	Adaptive	45.1	34.2 [†]	50.5 [†]

Table 3: Results of the user experiments. Static systems stand for conventional post-editing, without adaptation. Adaptive systems refer to post-editing in an environment with online learning. Time corresponds to the average post-editing time per sentence, in seconds. hTER and hBLEU refer to the TER and BLEU of the system hypothesis computed against the post-edited sentences. [†] indicates statistically significant differences between the static and the adaptive systems.

In terms of translation quality, adaptive systems performed much better than static ones, as reflected by the significant improvements in terms of hTER (5.3 and 4.2 points) and hBLEU (7.8 and 4.8 points). These results show that adaptive systems generated more correct translations, as they required less post-edits from the user.

In order to gain additional insights into the adaptation process, we studied the evolution of the hBLEU during the post-editing process. To this end, Fig. 2 compares the hBLEU per sentence of static and adaptive systems, for both test sets. Since the sentences were processed sequentially, we study the progress of the systems along its usage: for observing these trends, we computed a linear fit of the scores of each system via the least

squares method.

In Fig. 2a, we observe that for the test split T1, the adaptive system consistently produced slightly better hypotheses than the static one, but there was no clear evidence on the effects of online learning. Both systems behaved similarly: the hBLEU values were gradually increased, which suggests either that the test document was increasingly easy to translate or that the user felt more comfortable with the style and translations provided by the system. Therefore, they applied less post-edits to the final sentences.

In the case of T2 (Fig. 2b), we observe a degradation on the hBLEU of the static system, as the post-editing process advances. This degradation is prevented by the adaptive system, in which the hBLEU is even slightly increased. The effects of the adaptation are noticeable from the 30th sentence onwards.

Finally, it is interesting to compare the simulated experiment against this one. We observed that, in terms of automatic metrics, the system yielded much better results when evaluating against post-edits, rather than against reference sentences (compare the “Static” rows from Table 3 and Table 1, respectively). This suggests that the translation hypotheses provided by the system were useful to the human users, as they produced similar post-edited samples. It is also worth to point out that the adaptation process was, in most cases, slightly less effective in the simulated experiment.

5.3 User perceptions and opinions

After finishing each experiment, the participants answered a questionnaire regarding the post-editing task they had just performed. In this sur-

vey, we asked the users about their level of satisfaction of the translations they produced, whether they preferred to perform post-editing or translating from scratch and their opinions on the automatic translations provided, in terms of grammar, style and overall quality. We also requested for them to give their feedback on the task, as an open-answer question.

The users were generally satisfied with the translations they generated. In all cases, they preferred to perform this translation task via post-editing rather than translating from scratch. Two of them preferred to perform this translation from scratch in less than a 25% of the sentences. The other post-editor preferred to translate from scratch around a 50% of the sentences. In all cases, they are keen to perform translation post-editing in the future. These perceptions on the MT utility are slightly better than those reported by Daems and Macken (2019). We believe that these differences are due to the background in translation post-editing that our users had: they perform translation post-editing as their regular way of work; therefore, they perceptions toward this methodology are generally favorable.

Regarding the translation quality offered by the NMT system, their general opinion is that the system produced translations of average quality. The strongest attribute of the translations was their grammatical accuracy. The style and overall quality was perceived in some cases below the average, depending on the user and the experimental condition.

In order to avoid biases, the users did not know whether the experiment they performed featured a static or an adaptive system. Once they finished both experiments, they were asked to identify the adaptive systems. All users guessed correctly which one was the adaptive system.

Regarding their general opinions, they all observed how corrections applied on one segment were generally reflected in the following segment, especially corrections related to product names, grammatical structures and lexical aspects. This mostly reduced upcoming corrections to changes in the style. Overall, their perception was that the static system produced less fluent translations, and that the machine translation was very good in most cases, but useless in a few ones.

The post-editors reported a couple of minor issues regarding the NMT system: in a few cases,

they noticed that a domain-specific term was “forgotten” by the system, being wrongly translated. In addition, the users noticed in some cases, the occurrence of some made-up words (e.g., “absolvido”). This problem was probably caused by an incorrect segmentation of a word, via the byte pair encoding process. In order to deploy natural and effective translation systems, these problems need to be addressed.

6 Conclusions and future work

We conducted an evaluation of an adaptive NMT system in a post-editing scenario. The system leveraged the data generated during the post-editing process for adapting its underlying models. After testing the system in a laboratory setting, we conducted an experiment involving three professional translators, who regularly make use of MT post-editing. We observed reductions in post-editing times and significant improvements in terms of hTER and hBLEU, due to online learning. The users were pleased with the system. They noticed that corrections applied on a given segment generally were reflected on the successive ones, making the post-editing process more effective and less tedious.

As future work, we should address some of the concerns noticed by the post-editors, namely, the degradation of domain-specific terms and the incorrect generation of words due to subwords. To that end, we should study and analyze the hypotheses produced by the adaptive system and the post-edits performed by the users, similarly as Koponen et al. (2019). Moreover, we want to integrate our adaptive systems together with other translation tools, such as translation memories or terminological dictionaries, with the aim of fostering the productivity of the post-editing process. With this feature-rich system, we would like to conduct additional experiments involving more diverse languages and domains, using domain-specialized NMT systems, testing other models (e.g., Transformer, Vaswani et al., 2017) and involving a larger number of professional post-editors. Finally, we also intend to implement the interactive–predictive machine translation protocol (Lam et al., 2018; Peris and Casacuberta, 2019) in our translation environment, and compare it with the regular post-editing process.

Acknowledgements

The research leading to these results has received funding from the Spanish Centre for Technological and Industrial Development (Centro para el Desarrollo Tecnológico Industrial) (CDTI) and the European Union through Programa Operativo de Crecimiento Inteligente (Project IDI-20170964). We gratefully acknowledge the support of NVIDIA Corporation with the donation of a GPU used for part of this research, and the translators and project managers from Pangeanic for their help with the user study.

References

- Alabau, V., Bonk, R., Buck, C., Carl, M., Casacuberta, F., García-Martínez, M., González-Rubio, J., Koehn, P., Leiva, L. A., Mesa-Lao, B., Ortiz-Martínez, D., Saint-Amand, H., Sanchis-Trilles, G., and Tsoukala, C. (2013). CASMACAT: An open source workbench for advanced computer aided translation. *The Prague Bulletin of Mathematical Linguistics*, 100:101–112.
- Alabau, V., Carl, M., Casacuberta, F., García-Martínez, M., González-Rubio, J., Mesa-Lao, B., Ortiz-Martínez, D., Schaeffer, M., and Sanchis-Trilles, G. (2016). *New Directions in Empirical Translation Process Research*, chapter Learning Advanced Post-editing, pages 95–110. New Frontiers in Translation Studies.
- Arenas, A. G. (2008). Productivity and quality in the post-editing of outputs from translation memories and machine translation. *Localisation Focus*, 7(1):11–21.
- Aziz, W., Castilho, S., and Specia, L. (2012). Pet: a tool for post-editing and assessing machine translation. In *In proceedings of The International Conference on Language Resources and Evaluation*, pages 3982–3987.
- Bahdanau, D., Cho, K., and Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. *arXiv:1409.0473*.
- Bentivogli, L., Bertoldi, N., Cettolo, M., Federico, M., Negri, M., and Turchi, M. (2016). On the evaluation of adaptive machine translation for human post-editing. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 24(2):388–399.
- Biçici, E. and Yuret, D. (2015). Optimizing instance selection for statistical machine translation with feature decay algorithms. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 23(2):339–350.
- Bojar, O., Buck, C., Callison-Burch, C., Haddow, B., Koehn, P., Monz, C., Post, M., Saint-Amand, H., Soricut, R., and Specia, L., editors (2013). *Proceedings of the Eighth Workshop on Statistical Machine Translation*. Association for Computational Linguistics.
- Bojar, O., Haddow, B., , D. M., Sudarikov, R., Tamchyna, A., and Vari, D. (2017). Report on building translation systems for public health domain (deliverable D1.1). Technical Report H2020-ICT-2014-1-644402, Technical report, Health in my Language (HimL).
- Brown, P. F., Cocke, J., Pietra, S. A. D., Pietra, V. J. D., Jelinek, F., Lafferty, J. D., Mercer, R. L., and Roossin, P. S. (1990). A statistical approach to machine translation. *Computational Linguistics*, 16:79–85.
- Castilho, S., Moorkens, J., Gaspari, F., Calixto, I., Tinsley, J., and Way, A. (2017). Is neural machine translation the new state of the art? *The Prague Bulletin of Mathematical Linguistics*, 108(1):109–120.
- Chen, B. and Cherry, C. (2014). A systematic comparison of smoothing techniques for sentence-level bleu. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 362–367.
- Daems, J. and Macken, L. (2019). Interactive adaptive smt versus interactive adaptive nmt: a user experience evaluation. *Machine Translation*, pages 1–18.
- Denkowski, M., Dyer, C., and Lavie, A. (2014). Learning from post-editing: Online model adaptation for statistical machine translation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 395–404.
- Federico, M., Bertoldi, N., Cettolo, M., Negri, M., Turchi, M., Trombetti, M., Cattelan, A., Farina, A., Lupinetti, D., Martines, A., Massidda, A., Schwenk, H., Barrault, L., Blain, F., Koehn, P., Buck, C., and Germann, U. (2014). The matecat tool. In *Proceedings of the 25th International Conference on Computational Linguistics: System Demonstrations*, pages 129–132.
- Gehring, J., Auli, M., Grangier, D., Yarats, D., and

- Dauphin, Y. N. (2017). Convolutional sequence to sequence learning. *arXiv:1705.03122*.
- Gers, F. A., Schmidhuber, J., and Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural computation*, 12(10):2451–2471.
- Green, S., Heer, J., and Manning, C. D. (2013a). The efficacy of human post-editing for language translation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 439–448.
- Green, S., Wang, S., Cer, D., and Manning, C. D. (2013b). Fast and adaptive online training of feature-rich translation models. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 311–321.
- Hu, K. and Cadwell, P. (2016). A comparative study of post-editing guidelines. In *Proceedings of the 19th Annual Conference of the European Association for Machine Translation*, pages 34206–353.
- Jia, Y., Carl, M., and Wang, X. (2019). Post-editing neural machine translation versus phrase-based machine translation for english–chinese. *Machine Translation*, pages 1–21.
- Karimova, S., Simianer, P., and Riezler, S. (2018). A user-study on online adaptation of neural machine translation to human post-edits. *Machine Translation*, 32(4):309–324.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Klein, G., Kim, Y., Deng, Y., Senellart, J., and Rush, A. M. (2017). OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of the Association for the Computational Linguistics*, pages 67–72.
- Koponen, M., Salmi, L., and Nikulin, M. (2019). A product and process analysis of post-editor corrections on neural, statistical and rule-based machine translation output. *Machine Translation*, pages 1–30.
- Kothur, S. S. R., Knowles, R., and Koehn, P. (2018). Document-level adaptation for neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 64–73.
- Lam, T. K., Kreutzer, J., and Riezler, S. (2018). A reinforcement learning approach to interactive-predictive neural machine translation. In *Proceedings of the European Association for Machine Translation conference*, pages 169–178.
- Nielsen, J. (1993). *Usability Engineering*. Morgan Kaufmann Publishers Inc.
- Ortiz-Martínez, D. (2016). Online learning for statistical machine translation. *Computational Linguistics*, 42(1):121–161.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Peris, Á. and Casacuberta, F. (2019). Online learning for effort reduction in interactive neural machine translation. *Computer Speech & Language. In Press*.
- Peris, Á., Cebrián, L., and Casacuberta, F. (2017). Online learning for neural machine translation post-editing. *arXiv:1706.03196*.
- Post, M. (2018). A call for clarity in reporting bleu scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Riezler, S. and Maxwell, J. T. (2005). On some pitfalls in automatic evaluation and significance testing for mt. In *Proceedings of the ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 57–64.
- Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, 323(6088):533.
- Sennrich, R., Haddow, B., and Birch, A. (2016). Neural machine translation of rare words with subword units. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1715–1725.
- 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.

- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Proceedings of the Advances in Neural Information Processing Systems*, volume 27, pages 3104–3112.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9.
- Turchi, M., Negri, M., Farajian, M. A., and Federico, M. (2017). Continuous learning from human post-edits for neural machine translation. *The Prague Bulletin of Mathematical Linguistics*, 108(1):233–244.
- Vasconcellos, M. and León, M. (1985). SPANAM and ENGSPAN: machine translation at the pan american health organization. *Computational Linguistics*, 11(2-3).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, Ł., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv:1609.08144*.
- Wuebker, J., Simianer, P., and DeNero, J. (2018). Compact personalized models for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 881–886.