

# Post-editing Productivity with Neural Machine Translation: An Empirical Assessment of Speed and Quality in the Banking and Finance Domain

Samuel Lübli<sup>1,3</sup> Chantal Amrhein<sup>1,3</sup> Patrick Düggelein<sup>1</sup>

Beatriz Gonzalez<sup>2</sup> Alena Zwahlen<sup>1</sup> Martin Volk<sup>1,3</sup>

<sup>1</sup> TextShuttle AG    <sup>2</sup> Migros Bank AG

<sup>3</sup> Institute of Computational Linguistics, University of Zurich

## Abstract

Neural machine translation (NMT) has set new quality standards in automatic translation, yet its effect on post-editing productivity is still pending thorough investigation. We empirically test how the inclusion of NMT, in addition to domain-specific translation memories and termbases, impacts speed and quality in professional translation of financial texts. We find that even with language pairs that have received little attention in research settings and small amounts of in-domain data for system adaptation, NMT post-editing allows for substantial time savings and leads to equal or slightly better quality.

## 1 Introduction

The use of neural networks for sequence transduction (Kalchbrenner and Blunsom, 2013; Sutskever et al., 2014; Bahdanau et al., 2015) has led to astounding progress in the field of machine translation (MT), establishing a new level of quality in applications such as automatic news translation (Sennrich et al., 2016b; Hassan et al., 2018). Nevertheless, the creation of publication-grade translations still requires human involvement (Läubli et al., 2018; Toral et al., 2018a), and previous work has explored human-machine collaboration in the form of post-editing, where human experts revise machine translated text where necessary.

Empirical investigations of post-editing productivity with NMT are still scarce, especially for language combinations that do not include English as

either the source or target language. In this paper, we describe and discuss the results of a productivity test of domain-adapted NMT with the in-house translation team of Migros Bank, a mid-sized financial institution based in Zurich, Switzerland. We evaluate the use of NMT under typical working conditions, focussing on two rarely explored language pairs—German (DE) to French (FR) and Italian (IT)—and texts from a specialised domain: banking and finance. We show that using NMT in combination with translation memories (TMs) and terminology databases (TBs) enables professional translators to work faster with no loss (or slight gains) in quality, even with limited in-domain data for system adaptation.

## 2 Background

Early assessments of post-editing productivity were focussed on technical texts. While a study by Krings (1994) with user manuals for technical appliances and rule-based MT found mixed results, interest in post-editing grew with the advent of statistical MT (SMT), which enabled time savings of up to 40% in film subtitling (Volk, 2008; de Sousa et al., 2011) and software localisation (Plitt and Masselot, 2010). Subsequent work concluded that significant time savings can also be achieved in more complex domains such as legal (Federico et al., 2012) or marketing texts (Läubli et al., 2013).

Many productivity tests explored either translation from or into English (e.g., Plitt and Masselot, 2010), or translation between closely related languages such as Swedish and Danish (e.g., Volk, 2008). Green et al. (2013) conducted a large-scale experiment from English into three target languages with different canonical word order: Arabic (VSO), French (SVO), and German

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

Text	Topic	Source Words	Coverage FR				Coverage IT			
			100 %	90 %	80 %	<i>R</i>	100 %	90 %	80 %	<i>R</i>
A	Funding	1108	8.94	0.54	6.59	14.69	9.57	1.17	6.59	15.89
B	Funding	1006	2.58	3.68	4.17	9.23	1.29	1.69	1.39	3.93
C	Investing	1059	6.80	2.08	1.89	10.18	4.25	3.31	3.02	9.64
D	Investing	1077	14.48	10.58	0.84	24.68	7.24	15.51	3.06	23.65

**Table 1:** Source texts (DE) and their TM coverage in the target languages (FR, IT). Fuzzy bands include 90.00–99.99 % (referred to as 90 %) and 80.00–89.99 % (referred to as 80 %) matches. Coverage coefficients *R* indicate the percentage of translated words available in the TM for each text, considering 80–100 % matches.

(SOV). While post-editing was significantly faster than translation from scratch for all combinations, it is unclear whether their findings would equally apply to language pairs that do not include English, particularly if less MT training material is available. We investigate two language pairs that have received little attention in post-editing research: DE–FR and DE–IT.

The effect of using NMT rather than SMT on translation productivity has not yet been conclusively assessed. One of the first studies contrasting NMT and SMT quality found that NMT produces less morphological, lexical, and word order errors, thus reducing post-editing effort by 26 % in English to German subtitle translation (Bentivogli et al., 2016). However, post-editing effort was measured with HTER (Snover et al., 2006), a distance-based metric. Castilho et al. (2017) found that although more fluent, post-editing NMT rather than SMT output did not save time in an educational domain due to a higher number of omissions, additions, and mistranslations. Conversely, time savings doubled with NMT (+36 %) compared to SMT (+18 %) in literary translation (Toral et al., 2018b). The number of studies on NMT post-editing productivity is still limited, and further studies are needed, not least because findings obtained with different domains and language combinations are difficult to compare. The present study contributes data on NMT post-editing speed and quality in the financial domain.

Previous productivity tests used different experimental designs. In early work, Krings (1994) found that post-editing of rule-based MT resulted in a decrease in translation time by 7 % when translators used pen and paper, but an increase by 20 % when they used a computer instead. Plitt and Masselot (2010) and Green et al. (2013) compared post-editing to translation from scratch, using purpose-built web interfaces that showed one source sentence at a time, paired with a target text

box that was either populated with MT or empty. Proponents of field tests have argued that while improving control of extraneous variables, such designs reduce experimental validity in that they isolate translators from tools long indispensable in professional workflows, namely software workbenches that show multiple sentences at a time and suggestions from TMs and TBs (Federico et al., 2012; Läubli et al., 2013). We chose an in-situ design where translators had access to the tools and resources known from their daily work.

### 3 Assessment of Translation Productivity

We conducted a productivity test of domain-adapted NMT on the premises of Migros Bank. Subjects translated texts under two experimental conditions. In TM-ONLY, they used the translation workbench known from their their daily work, including a domain-specific TM, a domain-specific TB, and any online services (except machine translation) of choice. The same setup was used in POST-EDIT, except that sentences with no fuzzy match of at least 80 % in the TM were populated with MT within the translation workbench. We did not show MT where high fuzzy matches were available because editing high fuzzy matches is more efficient (Sánchez-Gijón et al., 2019).

**Materials** We used four German source texts from Migros Bank. The texts had not been translated by any of the translators involved in the experiment before, and had been excluded from the MT training material (see below). The TMs contained several exact and high fuzzy matches for each text (Table 1).

To pretranslate sentences in POST-EDIT, we trained WMT17-style bi-RNN systems (Sennrich et al., 2017) using the `marian` toolkit (Junczys-Dowmunt et al., 2016). The training material consisted of 6 million out-of-domain segments from publicly available OPUS corpora (Tiede-

Subject	Text	Seq.	MT	Words/h	Quality
FR-1	A	1	No	520.37	4.00
FR-1	B	2	No	630.82	5.50
FR-1	C	3	Yes	909.88	5.00
FR-1	D	4	Yes	602.56	5.00
FR-2	A	1	Yes	987.00	4.50
FR-2	B	2	Yes	1237.13	3.50
FR-2	C	3	No	682.64	4.00
FR-2	D	4	No	505.40	4.50
Average TM-ONLY			No	584.81	4.50
Average POST-EDIT			Yes	934.14	4.50
Difference (%)				59.74	0.00

(a) DE-FR

Subject	Text	Seq.	MT	Words/h	Quality
IT-1	A	1	No	389.41	4.00
IT-1	B	2	No	398.71	4.00
IT-1	C	3	Yes	647.87	4.50
IT-1	D	4	Yes	393.14	4.00
IT-2	A	1	Yes	401.19	5.50
IT-2	B	2	Yes	536.09	5.50
IT-2	C	3	No	553.00	5.50
IT-2	D	4	No	469.56	5.50
Average TM-ONLY			No	452.67	4.75
Average POST-EDIT			Yes	494.57	4.88
Difference (%)				9.26	0.13

(b) DE-IT

**Table 2:** Experimental conditions and results: the number of target words produced per hour (Words/h) and averaged overall impression scores (Quality) as assigned by two expert raters per translation.

mann, 2009), as well as 385'320 and 186'647 in-domain segments for FR and IT, respectively. We filtered both in- and out-of-domain segments through a set of mostly length-based heuristics (Zwahlen et al., 2016), and oversampled the former as a simple means of domain adaptation. While this has proven effective in other contexts (e. g., Sennrich et al., 2016a), we note that translation quality could likely be improved by means of more advanced techniques such as fine-tuning (Luong and Manning, 2015) or multi-domain modelling (Chu et al., 2017).

**Subjects** A total of four professional translators took part in the productivity test, two each for the target languages FR (FR-1, FR-2) and IT (IT-1, IT-2). All were members of Migros Bank's internal translation team. They were therefore familiar both with the software used and with the language and terminology of the documents to be translated. FR-1, who joined the organisation shortly before the experiment, was less experienced than the other participants. All subjects had been post-editing outputs of the MT systems used in the experiment (see above) for three months, and had received four hours of post-editing training.

**Procedure** Each subject translated the four German source texts in the same order. Conditions were counterbalanced (Table 2). Subjects were first briefed about the purpose and data collected during the experiment. They were then given 60 minutes to work on each text, which we announced would likely not be enough to translate all sentences. There were 10-minute breaks between working blocks, and a 30-minute break in the mid-

dle of the experiment. A post-experimental survey concluded the experiment.

We encountered no problems with data collection, with the exception of a temporary failure of IT-1's screen in the last working block. The device went into standby mode, which was not reported immediately and resulted in a total interruption of 4 minutes, which we deducted from the respective session before calculating translation speed as shown in Table 2.

### 3.1 Speed

We report translation speed as the number of target words produced per hour. To account for TM matches, we derive a TM coverage coefficient  $R$  for each text:

$$R = 1a + 0.9b + 0.8c, \quad (1)$$

where  $a$  is the percentage of 100 %,  $b$  the percentage of 90 %, and  $c$  the percentage of 80 % TM matches. We then adjust the number of words  $W$  translated in each experimental block as

$$W^* = (1 - R) W. \quad (2)$$

This approximation assumes uniform distribution of TM matches within texts.

Results are shown in Table 2. FR subjects produced 584.81 and 934.14 words per hour in TM-ONLY and POST-EDIT, respectively, an increase of 59.74 %. The difference was less marked in IT, with 452.67 and 494.57 words per hour produced in TM-ONLY and POST-EDIT, respectively (9.26 %).

Criterion	DE-FR		DE-IT	
	TM-ONLY	POST-EDIT	TM-ONLY	POST-EDIT
Coherence	4.75	5.25	5.00	5.00
Cohesion	4.75	4.50	5.25	5.00
Grammar	4.75	4.75	4.75	4.88
Style	4.50	5.00	5.00	5.00
Cultural adequacy	4.50	4.75	4.50	4.75
Overall Impression	4.50	4.50	4.75	4.88

**Table 3:** Detailed quality assessment results. Each cell is an average over eight scores: four translations scored by two expert raters. Overall impression was graded separately; it is not an average over the other criteria.

While focussing on descriptive statistics due to small sample size, we also fit linear-mixed effects models for inferential analysis. Carter and Wojton (2018) show that very small sample sizes can attain sufficient power when a single fixed effect factor is of interest, albeit at a greater risk of type I errors. We use experimental condition (TM-ONLY vs. POST-EDIT) as the fixed effect factor, and random intercepts for subjects and texts. The models show no deviation from homoscedasticity or normality in visual inspection of residual plots and Shapiro-Wilk tests. Likelihood ratio tests show a significant main effect of experimental condition in FR ( $\chi^2(1) = 9.74, p < .01$ ), but not IT ( $\chi^2(1) = 0.93, p = .33$ ).

### 3.2 Quality

The translations produced in the experiment were reviewed by university lecturers in professional translation, who were remunerated at standard hourly rates. Experts did not know which translations were produced using MT. The quality of each translation was independently assessed by two experts, who assigned scores on a 6-point scale (1 = worst, 6 = best) for coherence, cohesion, grammar, style, cultural adequacy, and overall impression.

Results are shown in Table 3. Each cell is an average over 8 scores: 4 texts evaluated by two experts. Note that experts assigned separate scores for overall impression, which may therefore deviate from the average over scores for the other criteria. Average per-text scores for overall impression are included in Table 2.

Considering overall impression, experts did not find a difference in quality between texts produced with and without MT in FR. In IT, texts translated with MT received slightly higher scores (+0.13). MT improved coherence (+0.50), style (+0.50), and cultural adequacy (+0.25) in FR, as well as

grammar (+0.13) and cultural adequacy (+0.25) in IT. Cohesion, on the other hand, was found to be better in texts produced without MT in both FR (-0.25) and IT (-0.25).

## 4 Discussion

While the minimum speed hardly differed between TM-ONLY and POST-EDIT, the latter allowed for higher average and maximum speed. In FR, the highest speed measured in POST-EDIT was 1237.13 words per hour (FR-2, text B), as opposed to 682.64 words per hour in TM-ONLY (FR-2, text C). In IT, the maximum speed in POST-EDIT was 647.87 words (IT-1, text C), and 553.00 words per hour in TM-ONLY (IT-2, text C).

Three out of four translators were faster in POST-EDIT on average. IT-2 did not benefit from MT: With an average speed of 511.28 words per hour in TM-ONLY and 468.64 words per hour in POST-EDIT, the subject was 8.34 % slower. Previous research has shown that not all translators benefit equally from MT (e. g., Plitt and Masselot, 2010; Koehn and Germann, 2014), which calls for large sample sizes in productivity tests (Green et al., 2013). Although improving robustness, involving a large number of translators is not always possible in practice – in our case, the in-house translation team had no more than four members, and involving external translators would have introduced other confounds (such as domain knowledge) that are hard to control for. We also note that IT-2 produced translations of above-average quality (Tables 2, 3), suggesting that MT may be less beneficial when aiming for maximum quality.

Another observation that warrants discussion is the difference in productivity between the two target languages. Again, one possible explanation is the small number of participants and measurements. A larger number of measurements would

allow more accurate conclusions to be drawn as to whether the maximum speed achieved in FR (FR-2, text B) is to be treated as an outlier, or if translators will repeatedly achieve a throughput of more than 1,000 words per hour with MT. Moreover, the DE–IT engine was trained with less in-domain material than the DE–FR engine. This resulted in lower raw MT quality for IT, which in turn may have resulted in lower productivity.<sup>1</sup> Screen recordings also showed that IT translators made more stylistic changes to MT outputs, but apart from slightly higher quality scores overall (Table 3), we cannot quantify this finding and leave a more detailed analysis to future work.

With respect to quality, our results confirm previous findings that post-editing leads to similar or better translations (e. g., Green et al., 2013). An interesting nuance is that we find a slight, but consistent decrease in textual coherence within post-edited translations in both language pairs. As the research community is increasingly focussing on document-level MT, translation workbench providers will need to ensure integrability for future experimentation in real-life settings.

## 5 Conclusion

We have assessed the impact of NMT on translation speed and quality in the banking and finance domain. Despite working with language pairs that have received limited attention in research contexts and employing a simple means of domain adaptation, the use of NMT enabled professional translators to work faster: 59.74 % in DE–FR and 9.26 % in DE–IT. Unlike a number of previous studies, these improvements are not relative to translation from scratch, but to translation with domain-specific TMs and TBs within a customary translation workbench, which sets a higher baseline in terms of translation speed.

NMT did not have a negative impact on quality. To the contrary, scores assigned by expert raters were slightly higher for post-edited DE–IT translations. Screen recordings showed that IT translators devoted more time to stylistic changes of NMT output, underpinning the importance of translator training in cases where NMT is to be used to optimise throughput rather than quality.

Another factor that likely contributed to the dif-

ference between time savings in DE–FR and DE–IT is that roughly half as much in-domain training data was available for the latter. While further investigation will be needed to determine the impact of in-domain data volume and more advanced domain adaptation techniques, our results suggest that NMT has the potential of increasing translation productivity even with complex text types, little-researched language pairs, and limited amounts of in-domain training data. The present study contributes empirical evidence for DE–FR and DE–IT translation of financial texts, and we hope to encourage similar investigations with other languages and domains.

## References

- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.
- Bentivogli, Luisa, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. Neural versus phrase-based machine translation quality: a case study. In *Proceedings of EMNLP*. Austin TX, USA, pages 257–267.
- Carter, Kristina A. and Heather M. Wojton. 2018. The effect of extremes in small sample size on simple mixed models: A comparison of level-1 and level-2 size. Technical Report NS D-8965, IDA, Alexandria VA, USA.
- Castilho, Sheila, Joss Moorkens, Federico Gaspari, Iacer Calixto, John Tinsley, and Andy Way. 2017. Is Neural Machine Translation the New State of the Art? *The Prague Bulletin of Mathematical Linguistics* 108:109–120.
- Chu, Chenhui, Raj Dabre, and Sadao Kurohashi. 2017. An empirical comparison of domain adaptation methods for neural machine translation. In *Proceedings of ACL*. Vancouver, Canada, pages 385–391.
- de Sousa, Sheila C. M., Wilker Aziz, and Lucia Specia. 2011. Assessing the post-editing effort for automatic and semi-automatic translations of dvd subtitles. In *Proceedings of RANLP*. Hissar, Bulgaria, pages 97–103.
- Federico, Marcello, Alessandro Cattelan, and Marco Trombetti. 2012. Measuring user productivity in machine translation enhanced computer assisted translation. In *Proceedings of AMTA*. San Diego CA, USA.
- Green, Spence, Jeffrey Heer, and Christopher D. Manning. 2013. The efficacy of human post-editing for language translation. In *Proceedings of CHI*. Paris, France.
- Hassan, Hany, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming

<sup>1</sup>However, Koehn and Germann (2014) find that between-subjects variance is higher than between-systems variance in post-editing.

- Zhou. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint* 1803.05567.
- Junczys-Dowmunt, Marcin, Tomasz Dwojak, and Hieu Hoang. 2016. Is neural machine translation ready for deployment? A case study on 30 translation directions. In *Proceedings of IWSLT*. Tokyo, Japan.
- Kalchbrenner, Nal and Phil Blunsom. 2013. Recurrent continuous translation models. In *Proceedings of EMNLP*. Seattle, Washington, USA, pages 1700–1709.
- Koehn, Philipp and Ulrich Germann. 2014. The impact of machine translation quality on human post-editing. In *Proceedings of the EACL 2014 Workshop on Humans and Computer-assisted Translation*. Gothenburg, Sweden, pages 38–46.
- Krings, Hans P. 1994. *Texte reparieren: Empirische Untersuchungen zum Prozeß der Nachredaktion von Maschinenübersetzungen*. Habilitation thesis, Universität Hildesheim, Hildesheim, Germany.
- Luong, Minh-Thang and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domain. In *Proceedings of IWSLT*. Da Nang, Vietnam.
- Läubli, Samuel, Mark Fishel, Gary Massey, Maureen Ehrensberger-Dow, and Martin Volk. 2013. Assessing post-editing efficiency in a realistic translation environment. In *Proceedings of the 2nd Workshop on Post-Editing Technology and Practice (WPTP)*. Nice, France, pages 83–91.
- Läubli, Samuel, Rico Sennrich, and Martin Volk. 2018. Has machine translation achieved human parity? A case for document-level evaluation. In *Proceedings of EMNLP*. Association for Computational Linguistics, Brussels, Belgium, pages 4791–4796.
- Plitt, Mirko and François Masselot. 2010. A productivity test of statistical machine translation post-editing in a typical localisation context. *Prague Bulletin of Mathematical Linguistics* 93:7–16.
- Sennrich, Rico, Alexandra Birch, Anna Currey, Ulrich Germann, Barry Haddow, Kenneth Heafield, Antonio Valerio Miceli Barone, and Philip Williams. 2017. The University of Edinburgh’s neural MT systems for WMT17. In *Proceedings of WMT*. Copenhagen, Denmark.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016a. Edinburgh neural machine translation systems for WMT 16. In *Proceedings of WMT*. Berlin, Germany, pages 371–376.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016b. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of ACL*. Berlin, Germany, pages 1715–1725.
- Snover, Matthew, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of AMTA*. Boston MA, USA.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*. Montreal, Canada, pages 3104–3112.
- Sánchez-Gijón, Pilar, Joss Moorkens, and Andy Way. 2019. Post-editing neural machine translation versus translation memory segments. *Machine Translation* (4 April 2019).
- Tiedemann, Jörg. 2009. News from OPUS – A collection of multilingual parallel corpora with tools and interfaces. In *Proceedings of RANLP*. Borovets, Bulgaria, pages 237–248.
- Toral, Antonio, Sheila Castilho, Ke Hu, and Andy Way. 2018a. Attaining the unattainable? Reassessing claims of human parity in neural machine translation. In *Proceedings of WMT*. Association for Computational Linguistics, Belgium, Brussels, pages 113–123.
- Toral, Antonio, Martijn Wieling, and Andy Way. 2018b. Post-editing effort of a novel with statistical and neural machine translation. *Frontiers in Digital Humanities* 5(9).
- Volk, Martin. 2008. The automatic translation of film subtitles: a machine translation success story? In J Nivre, M Dahllöf, and B Megyesi, editors, *Resourceful Language Technology: Festschrift in Honor of Anna Sågvall Hein*, Uppsala University, Uppsala, Sweden, Studia Linguistica Upsaliensia, pages 202–214.
- Zwahlen, Alena, Olivier Carnal, and Samuel Läubli. 2016. Automatic TM cleaning through MT and POS tagging: Autodesk’s submission to the NLP4TM 2016 shared task. *arXiv preprint* 1605.05906.