Nunes Vieira, L. (2019). Post-Editing of Machine Translation. In M. O'Hagan (Ed.), *The Routledge Handbook of Translation and Technology* (1 ed., pp. 319-335). Routledge.
<https://doi.org/10.4324/9781315311258>

Peer reviewed version

License (if available):
Unspecified

Link to published version (if available):
[10.4324/9781315311258](https://doi.org/10.4324/9781315311258)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via Routledge at <https://doi.org/10.4324/9781315311258>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Vieira, L.N. (2019). Post-Editing of Machine Translation. In M. O'Hagan (Ed.), *The Routledge Handbook of Translation and Technology*. Routledge (pp. 319-335)

Link to published book: <http://www.routledge.com/9781138232846>

Post-Editing of Machine Translation

Lucas Nunes Vieira

Abstract

This chapter analyses the evolutionary process that human post-editing of Machine Translation (MT) output has undergone in previous years. It is well known that post-editing can improve translating productivity as well as target-text quality relative to translation carried out 'from scratch', but post-editing has many facets. Initially, it was akin to a step in the MT development pipeline where humans tidied the machine output to make it usable. More recently, post-editing is a professional service with its own international standard. In addition, the focus of translation tasks involving MT is shifting towards more integrated and human-centred environments. The chapter reviews previous post-editing research and discusses emerging issues and future directions on the use of MT in professional translation. It outlines how post-editing has matured as a practice and service and adopts an encompassing understanding of post-editing tasks that reflects MT's changing role in professional translation processes.

Keywords

Post-editing, interactive post-editing, static post-editing, machine translation, language services, translators, post-editors, agency

Introduction

In most professional contexts, the output of MT systems needs to be post-edited to reach the desired standards of quality. Post-editing of MT has been carried out since the early days of MT technology, but it has evolved considerably in the last few years as a practice, service and research topic. At the inception of MT research, while it was acknowledged that interactions between humans and MT systems were inevitable for translations of high-stakes content, MT researchers saw post-editing as an undesirable 'final step' in MT development (Reifler 1952: n.p.). Some of the language used to describe post-editing at the time clearly alludes to a passive activity where human editors helped to close the gap between defective MT outputs and MT's ultimate purpose of providing fully automatic high-quality translations (see Läubli and Green; Melby in this volume). Post-editors were referred to as MT's 'human partners' (Bar-Hillel 1951: 230). In some cases, they were not required to know the source language. It was thought that by combining non-expert source-language pre-editors – i.e. those in charge of making the source text more MT-friendly (see 'Effort Prediction', below) – with non-expert post-editors it would be possible to increase translation provision by reducing the need for scarce bilingual translators (*ibid.*). Post-editing therefore came about as part of a paradigm where human editors assisted the machine rather than one where the machine assisted them. Given this machine-centred nature of the activity and the poor results – from today's perspective – that MT was able to deliver then, it is not surprising that post-editing has

developed an unfavourable reputation throughout the history of MT. While there are reports of translators' acceptance of MT (e.g. Guerberof 2013), translators' attitudes to MT and post-editing are often negative (Moorkens and O'Brien 2015, Läubli and Orrego-Carmona 2017).

However, much has been done lately to improve the use of MT in professional translation. MT is now available in most computer-assisted translation (CAT) tools. As implied by their name, these tools changed the focus from human-assisted MT to machine- (or computer-) assisted human translation. Unlike the paradigm from early MT research, machine-assisted human translation puts humans at the centre of translation production. The incorporation of MT into CAT tools and, more recently, a higher degree of integration of these tools' different features – such as translation memories, terminological resources and MT itself – all contribute to an understanding of MT as a resource that helps human translators improve their productivity. On the other hand, integrating different technologies into CAT environments has also blurred the lines between technologies and different sources of assistance that can be used in the translation process. Consequently, post-editing is now in a state of terminological flux where it can be seen to comprise different tasks and procedures.

This chapter discusses the current state of post-editing and the evolutionary process it has undergone. It adopts an encompassing approach to the subject that sees post-editing not only as a separate service, which now has its own international standard (ISO 2017), but also as a multifaceted CAT activity that goes beyond just static tidying of MT outputs. The remainder of the chapter is structured as follows. First, findings from previous research are discussed in relation to two important factors: quality and post-editing effort. Second, emerging issues are analysed with respect to the use of Neural MT (NMT) in post-editing tasks and the post-editing of literary content. Third, post-editing is examined through the lens of human agency and the different ways in which post-editing as a service and activity can currently be understood. Finally, the chapter concludes with a summary and an overview of potential future directions regarding the role of MT post-editing in professional translation and translation studies.

Effort and Quality in Post-Editing

Most research on post-editing to date revolves around the key concepts of quality and effort. For one thing, the potential productivity improvements provided by MT would be undermined if the quality of post-edited texts were to be considered sub-standard. Moreover, if raw MT quality is low, post-editing may require too much effort, which would not make it worthwhile relative to translating the source text from scratch. Therefore, effort and quality – be it MT quality or target-text quality – are important factors in determining post-editing's feasibility and the benefits of different post-editing levels and modalities. These two concepts underpin the discussion presented in this section. The discussion first focuses on two perspectives from which post-editing tasks can be classified: a source-text-exposure perspective where post-editing can be presented as bilingual or monolingual tasks – i.e. where both the source text and the MT output are available or where only the MT output is accessed during editing, respectively – and an editing mode perspective, where post-editing can be static or interactive. Subsequently, the section discusses effort prediction, the post-edited product and post-editing guidelines. The chapter's focus is on human post-editing, so automatic post-editing is not covered in detail here. For introductory comments on automatic post-editing, see the 'Agency in the post-editing process' section.

Bilingual and Monolingual Post-Editing

In most instances, including in the present chapter, ‘post-editing’ is assumed to mean ‘bilingual post-editing’. This is consistent with how post-editing is carried out in commercial translation tools where both the source and the target text are displayed on screen by default. The feasibility of monolingual post-editing has nevertheless also been an object of research. The rationale for research on this issue tends to be that, if monolingual post-editing were indeed shown to be worthwhile, it would cut costs by simplifying the task and potentially reducing the need for bilingual expertise. A first observation to make in this respect is that monolingual post-editing is usually more promising where the MT output is edited by domain specialists. A previous study reports that over 90% of sentences post-edited monolingually by a domain expert were found to be completely correct (Schwartz 2014: 34). However, other research on this subject has to date only partially confirmed monolingual post-editing to be viable. Studies comparing monolingual and bilingual conditions usually show that monolingual post-editing improves the MT output in terms of fluency (i.e. intrinsic linguistic quality) but that in terms of adequacy – or the extent to which the source-text meaning is conveyed – bilingual post-editing is unsurprisingly superior (Mitchell, Roturier, and O’Brien 2013, Nitzke 2016).

Given the inherent risk of ignoring the source text, in professional settings – including those covered by the ISO 18587 standard – checking the MT output against the source-language content is usually a strict requirement, so bilingual post-editing is currently the norm. To examine the feasibility of incorporating MT into human translation, a vast amount of previous research on bilingual post-editing has compared post-editing to from-scratch translation in terms of effort and productivity. A relatively early study to undertake this comparison is also the study that provided a now recurrent framework for measuring the level of effort post-editing requires (Krings 2001). In Krings’s formulation, overall post-editing effort can be divided into three effort types: cognitive (i.e. the mental processing behind editing decisions), technical (i.e. the mechanical effort involved in implementing the edits) and temporal (i.e. the time required by the activity, which comprises both cognitive and technical effort) (*ibid.*: 178-179). These three effort dimensions are often the basis for comparing post-editing to from-scratch translation.

The general pattern in previous research investigating post-editing’s viability is that post-editing requires less effort relative to from-scratch translation and that it does not have a detrimental effect on target-text quality. However, previous studies report some variation in how useful MT is found to be. This may be down to the language pair or the fact that certain texts may be more conducive to successful post-editing (see ‘Effort Prediction’, below). While an exhaustive analysis of these results is beyond the scope of this chapter,ⁱ two studies are mentioned here for the diversity of language pairs they analyse and for the fact that their participant samples consist of professional translators. The first study is based on machine translations of Wikipedia content from English into Arabic, French and German (Green, Heer, and Manning 2013). This study reported an increase in productivity for the post-editing condition in all these language directions. It also reported higher target-text quality for the post-editing condition across the languages analysed, which suggests that slower from-scratch translation is not necessarily more effective. The second study is a productivity test carried out at Autodesk (Zhechev 2014). This study tested translators’ productivity for post-editing and for translation from scratch based on tasks carried out from English into nine target languages: French, Korean, Italian, Brazilian Portuguese, Spanish, Japanese, Simplified Chinese, German and Polish. The study reported productivity increases for all these languages ranging between 37.13% for Polish and 92.33% for French (*ibid.*: 9). While this study provided no details of post-edited quality, in a previous productivity test conducted at Autodesk (Plitt and Masselot 2010), the post-edited content was found to contain fewer errors compared to content translated from scratch, which is consistent with results reported by Green *et al.* (2013).

Studies like the ones mentioned above have helped to demonstrate that post-editing can be an attractive activity. Subject to appropriate conditions concerning text suitability and robust – preferably in-domain – MT training, the potential productivity benefits of post-editing are now well established. This field nevertheless has several open issues concerning, not least, how post-editing tasks should be carried out. Some of these issues are discussed below.

Static and Interactive Post-Editing

The studies reviewed so far are based largely on the static post-editing paradigm. In this paradigm, the MT output is generated first and then it is edited statically as a separate step. A different paradigm for MT use in professional translation is one where translators interact with MT systems while the final version of the text is generated. This may happen in situations where MT can be used to predict and complete the human translations as they are typed, or it may happen in a reciprocal interactive manner, where the MT system reacts to and learns from the human edits on the fly. A commercial translation tool where this second interactive mode is implemented by default is Liltⁱⁱ, a tool that also integrates translation memories and term bases by using them to fine-tune the machine suggestions. Interactive editing has the potential to change the post-editing process and its product, so previous research has examined how interactive conditions affect translators' work. A full discussion of previous research on interactive MT systems is provided in Läubli and Green in this volume. Results that concern translators' work and post-editing processes and products more directly are reviewed below.

Previous investigations of interactive post-editing with a focus on professional translation showed mixed results. Initial user evaluation studies of the tool TransType showed that the interactive condition, which then consisted mostly of using MT for purposes akin to predictive typing, led to decreases in productivity of up to 35% relative to the static editing condition (Langlais and Lapalme 2002: 90). Different types of interactive models – e.g. with online learning, where the MT system learns on the fly from translators' edits – were tested during the field trials of the CASMACAT project. Results from these field trials suggested that interactive conditions required less technical effort (i.e. fewer keystrokes) (Koehn *et al.* 2015: 23-28). They also confirmed a previous assumption that post-editors gradually become faster as they get used to the interactive mode, a human learning effect that could not be observed for static post-editing (*ibid.*). Based on a snapshot (i.e. non-longitudinal) analysis comparing static and interactive editing conditions, a separate study also conducted in the CASMACAT workbench reported more editing time and higher technical effort for interactive post-editing, though shorter eye fixations on screen, a sign of less cognitive effort (Alves *et al.* 2016).

The impact of interactive editing on the quality of post-edited products has been examined less frequently than its impact on productivity. Some previous research in this respect suggests that interactive post-editing takes more time but that it may lead to products of higher quality relative to static post-editing (e.g. Green *et al.* 2014). Conversely, it has also been shown that interactive post-editing does not have an appreciable effect on target-text quality and that on occasion it may even make editors add errors to the text (Underwood *et al.* 2014: 557). It is worth noting that the study by Underwood *et al.* is one of few that rely on human quality evaluations in examining the efficacy of interactive conditions (see also Koehn and Haddow 2009). Other studies rely mostly on automatic assessment methods, so given current tendencies towards higher interactivity in translation production, the use of robust human evaluation methodologies to examine this subject is an issue that probably calls for further research (see also Doherty in this volume).

Effort Prediction

Post-editing's benefits notwithstanding, factors such as source-text genre and complexity are known to have an impact on MT quality, which in turn can affect the level of effort translators need to expend on post-editing tasks. There is a strand of research that therefore attempts to predict how much effort post-editing is likely to require – or indeed if it is worth at all – by using source-text and/or MT-output characteristics as potential predictors. Exploiting source-text characteristics in this context is linked to the concept of pre-editing, which consists of ridding the source text of features that are known to pose a challenge for MT (see e.g. Bernth and Gdaniec 2002). More recently, a series of studies has attempted to identify these features based on characteristics such as sentence length, syntactic structure, the count and distribution of different parts of speech as well as automatic MT evaluation scores (e.g. Vieira 2014, Green, Heer and Manning 2013, O'Brien 2011, Tatsumi and Roturier 2010, Aziz, Koponen, and Specia 2014). Most of these studies are based on English source texts, where the incidence of nouns has been reported as a potential effort predictor (see Green *et al.* 2013; Aziz *et al.* 2014). Research based on French source texts also suggests the incidence of prepositional phrases as an effort-predicting feature (Vieira 2014). Regarding sentence length, it has been found that longer sentences pose more cognitive effort in post-editing (Koponen 2012), but that it should not be taken as a given that sentence length will always correlate strongly with post-editing time (Aziz *et al.* 2014).

While predicting effort is a research topic in translation studies more generally, effort predictors in an MT context can have specific instrumental purposes. These predictors can be exploited to estimate MT quality automatically without the need for human reference translations – an area of research referred to as quality estimation (see Specia 2011). This technology was examined in the CSMACAT project, where it was surprisingly the least popular feature among translators who took part in the project's field trials (Koehn *et al.* 2015: 25). However, research on this subject is still evolving, for example in the context of annual shared tasks organised within the Workshop on Machine Translation (Specia *et al.* 2018). Recent interviews with professional translators (Vieira and Alonso 2018) suggest that being able to see problematic MT passages in the same way that most CAT tools mark suggested edits in translation memory matches would be a welcome development. As quality estimation technology improves and gets closer to reaching desirable levels of usefulness, we may see it being applied to features of this kind in mainstream CAT tools in the future.

The Post-Edited Product

Research on post-editing usually measures the quality of the post-edited product as a complement to measurements of productivity when comparing post-editing to from-scratch translation. However, interest in the post-edited product has increased lately and this subject is now examined from different angles as an issue in its own right. An increasingly popular method for investigating post-edited quality is the involvement of end-users in the research. As MT technology improves, a relevant question to ask is whether there is a difference between raw MT outputs and corresponding post-edited versions in how well these texts are found to fulfil their real-world functions. A previous study investigated this based on software instructions that were machine-translated from English into Brazilian Portuguese (Castilho *et al.* 2014). Even by restricting edits just to the ones necessary to ensure accuracy and comprehensibility, this study found that post-editing significantly increased MT's usability based on measures including eye tracking metrics (see Jakobsen in this volume) and self-reported satisfaction.

Similarly, a recent study based on Welsh carried out a reading task involving eye tracking and self-reported scores of comprehensibility and readability (Screen 2019). This study compared post-edited

translations to translations produced from scratch. While post-edited translations were not found to be superior, this study found the two types of product to be largely equivalent, which like previous research supports the use of MT in professional settings. Another recent study examines how different levels of editing (see ‘Post-Editing Levels and Guidelines’, below) affect the reception of post-edited texts (van Egdom and Pluymaekers 2019). This study tested four levels ranging between ‘minimal’ and ‘full’ post-editing. Especially for informative texts, this study found the highest post-editing level to be largely redundant as in most cases this level did not improve end-users’ perception of the content compared to ‘moderate’ post-editing, the third level (ibid.: 65). These findings are consistent with previous work suggesting that high levels of editing and of cognitive effort do not necessarily increase post-edited quality in a linear fashion (Vieira 2017).

Post-Editing Levels and Guidelines

Documents that are not intended to be published openly for dissemination purposes may well tolerate stylistic issues that do not affect comprehensibility so, especially in static post-editing tasks, the MT output can be edited more or less drastically depending on the text’s real-world context of use and its requirements. This notion of fit-for-purpose translation (see Bowker in this volume) and of how the text’s intended use should dictate the task requirements is not new in translation (see Vermeer 2000), but MT brings a new perspective into this debate in that it allows texts that are intrinsically problematic to be put into use so that specific goals are achieved.

The notions described above are directly linked to different levels of post-editing proposed in guidelines published in industrial settings. The number and description of these levels can vary quite considerably. An early report of how post-editing was carried at the European Commission mentions a ‘rapid’ and a ‘full’ post-editing level, where the main differences between the two were the time spent on the task and the quality of the final product (Wagner 1985). ‘Minimal’ post-editing has also been mentioned as a fuzzy level in between ‘rapid’ and ‘full’ (Allen 2003), and there have also been studies like the one by van Egdom and Pluymaekers (2019), where four levels of post-editing are adopted – in van Egdom and Pluymaekers’s case, ‘minimal’, ‘light’, ‘moderate’ and ‘full’, specifically.

Even though approaches to post-editing levels are wide-ranging, the most popular levels are often referred to as just ‘light’ and ‘full’ post-editing (see Hu and Cadwell 2016). Influential guidelines published by the Translation Automation User Society (TAUS) refer to these levels based on two standards of expected target-text quality, namely ‘good enough’ quality and quality which is ‘similar or equal to human translation’ (Massardo *et al.* 2016: 17-18). These two quality standards roughly correspond to ‘light’ and ‘full’ post-editing, respectively. However, by describing levels of post-editing from the perspective of expected target quality, the TAUS guidelines reflect the fact that the actual amount of editing to be carried out inevitably depends on the text’s intended purpose as well as on the quality of the raw MT output. This means that target quality is likely to be a more consistent parameter for post-editing guidelines than the actual degree of editing to be carried out.

Instructions regarding the different issues post-editors should focus on when working under the ‘good enough’ and the ‘human translation quality’ levels are nevertheless provided in the TAUS guidelines. These instructions vary mostly with respect to whether MT errors affect meaning. Specifically, when working towards ‘good enough’ quality, semantics and comprehensibility should be the focus, rather than syntactic or grammatical issues that do not affect comprehension. When working towards human translation quality, however, issues with style, syntax, grammar and formatting should also be addressed, as should terms that need to remain in the original language but which may have been translated by the MT system (Massardo *et al.* 2016: 17-18). Requirements for full post-editing published in the ISO 18587 standard largely follow the TAUS guidelines. The ISO

18587 also provides guidelines for light post-editing that similarly follow the TAUS documentation, but this is done for informative purposes only; light post-editing is not within the remit of the ISO 18587.

While a comprehensive review of different post-editing guidelines cannot be undertaken here, the study by Hu and Cadwell (2016) concludes that full post-editing is where recommendations vary the most, especially with regard to the requirement for stylistic edits. In interviews conducted by the present author with translators and translation companies, considerable variation was observed in how these guidelines are put into practice (Vieira and Alonso 2018: 4-5). Some companies regarded post-editing as a task that was by default associated with lower quality expectations, a context where the distinction between 'light' and 'full' post-editing is less relevant (*ibid.*). There were also companies that closely followed the TAUS guidelines and expected translators to adjust the level of editing depending on the commission, but this was often found to be problematic either because translators found it hard to attain themselves strictly to one of these levels or because the 'light' edit carried out was not different enough from 'full' post-editing (*ibid.*). Given the shift towards human-centred tasks where MT is regarded as a tool (see 'Agency in the post-editing process', below), the notion of post-editing levels may well be phased out or give way to a different concept in the future, which is discussed in the Conclusion.

Emerging Issues

Post-Editing of Neural MT

The emergence of Neural MT (NMT) (see Melby in this volume) has the potential to change certain aspects of how post-editing tasks are carried out. While at the time of writing this MT architecture is broadly regarded as the state of the art in MT technology, previous research has shown that phrase-based statistical machine translation (PBSMT) may have superior performance in transmitting source-text meaning despite the higher fluency of NMT outputs (Castilho *et al.* 2017). For one thing, this means that NMT could make monolingual post-editing (see 'Bilingual and Monolingual Post-editing', above) even more challenging since adequacy errors can be difficult – or indeed impossible – to address without access to the source text, especially for fluent translations.

The higher fluency of neural systems may have the potential of making MT errors harder to identify and correct even in bilingual editing conditions. A recent study comparing post-editing of PBSMT with post-editing of NMT found that translation students had a poorer error correction rate when post-editing NMT despite the fact that the product of their work in the NMT condition was found to be superior (Yamada 2019: 87). The NMT output used in this study, which was based on English-to-Japanese tasks, had fewer errors compared to the PBSMT output. However, there was no significant difference in the levels of cognitive effort expended by the students under the two conditions, which suggests that NMT made the students concentrate higher levels of effort on fewer edits. This result led Yamada (*ibid.*) to conclude that NMT is unlikely to enhance the performance of translators who are still in training. Similarly, he posits that the rise of NMT makes it even more important that professional translators, rather than non-professionals, carry out post-editing tasks (see King in this volume).

Moreover, research comparing post-editing of NMT to from-scratch translation has observed a text genre effect in tasks carried out from English to Simplified Chinese (Jia, Carl and Wang 2019: 60). In this study, cognitive effort was overall lower for NMT post-editing. However, while no significant difference in target-text quality was observed between from-scratch translation and NMT post-

editing, the NMT condition was only found to be faster for texts of a specialized nature, namely a patient information leaflet and a dishwasher manual (*ibid.*). While participants in this study were again trainee translators, the fact that NMT post-editing was not faster for unspecialized texts shows that, despite recent advances, MT technology still has much to improve. Given commercial MT developers' decades-old tendency to oversell the results of their systems – see Hassan *et al.* (2018) and a rebuttal by Toral *et al.* (2018) – translation studies have an important role to play in ensuring a sound understanding of MT's place in professional translation. This is likely to become even more critical as MT technology continues to evolve and gain wider public visibility.

Post-Editing of Literary Content

An emerging research area is the use of NMT for post-editing of literary content. A recent study has found that using NMT to post-edit a novel led to a 36% increase in translators' words-per-time productivity relative to translation from scratch (Toral, Wieling and Way 2018). Surprisingly, this same study found that even PBSMT led to an 18% increase in translators' productivity compared to the from-scratch condition. The MT system used in this study was tailored to literary content (Toral and Way 2018), which is likely to have played a significant role in translators' productivity improvements. At the time of writing, target-text quality evaluations are not available for the post-edited texts produced in this experiment. In any case, in a previous study based on the PBSMT paradigm, a post-edited literary text was found to be 'acceptable' according to the assessment of language teaching assistants and a professional translator (Besacier 2014). While the professional translator playing the role of assessor in Besacier's study had mixed feelings about the post-edited text's quality (*ibid.*: n.p.), these results are a promising sign that literary translation too may be an area where MT can be a useful tool.

Investigations in this area are on the rise, but unsurprisingly literary translators do not prefer post-editing over from-scratch translation (see Moorkens *et al.* 2018). In addition, the use of MT for literary post-editing brings about several issues. First, the very concept of words-per-time productivity may not have as much value for this type of content as it does for texts in technical domains. This is because in literary translation – perhaps more so than in other areas of specialization – speed may be less critical than factors such as creativity, aesthetics and the readers' experience. Depending on the context, target quality is therefore an even more complex variable for post-editing research involving literary texts. This, in turn, poses other issues. The emphasis on traditional assessment methodologies employed in MT research, such as automatic evaluations and sentence-level assessments, may need to be re-examined if research in this area is to stay true to readers' and literary translators' perceptions of quality (see Doherty in this volume). The BLEU metric, for instance (Papineni *et al.* 2002) – the workhorse of automatic evaluations among MT developers whose variations are sometimes used to evaluate post-edited content (see Green *et al.* 2014) – would need to give way to reader-centred evaluation approaches. This is because the concept of reference translations, the basis for generating BLEU scores, can easily be considered unfit for purpose in literary contexts, where overlapping versions of the same source text can differ quite substantially and still be considered acceptable. This is arguably the case for many other types of translation, but in literary domains these issues merit special attention. Regarding the evaluation granularity, while recent research shows the value of document-level assessments (Läubli, Sennrich and Volk 2018), sentence-level assessments are still the norm. Especially in the context of literary post-editing, document-level evaluation is argued here as a preferable future direction. This, however, poses issues of its own as the higher the granularity level the more complex (and expensive) evaluation tasks arguably become. In literary post-editing research, it may be practically

hard to capture readers' experience in relation to the totality of a novel or even a longer short story, for instance.

Agency in the post-editing process

The discussion presented so far has touched upon various facets of post-editing, including static and interactive editing modes. Given the diverse ways in which MT can be used in translation production, as mentioned in the Introduction the term 'post-editing' can currently be seen to refer to different tasks. Instability in what post-editing is seen to involve lies mainly in the use of MT as an optional source of translation suggestions in CAT tools or where MT is used interactively in the translation process. MT being used in these ways may in some cases be classed as something other than post-editing. The translation tool Lilt refers to this MT use paradigm as 'interactive and adaptive machine suggestions', which is described in the tool's documentation as a replacement for post-editingⁱⁱⁱ. It has also been proposed that a paradigm where MT and other technologies assist translators in an integrated production framework could be classed as 'augmented translation', which can involve automating even certain aspects of project management (Lommel 2018).

Interacting with MT in these contexts, in one way or another and to different extents, still involves having human translators use and edit MT output, so the argument that completely new activities have replaced post-editing is debatable. Nevertheless, while there can be many drivers for avoiding the term post-editing in these situations – not least translators' negative attitudes to the activity – the fact that translators may acquire a higher level of control over MT use as the paradigm shifts from human-assisted MT to machine-assisted human translation cannot be ignored. Therefore, a key factor in the currently diverse nature of post-editing, it is proposed here, concerns the extent to which human translators are expected to exercise agency in the post-editing process. The concept of agency can be understood as the 'willingness and ability to act' (Kinnunen and Koskinen 2010: 6). While a comprehensive discussion of this concept cannot be provided here, this section draws on this notion mostly with respect to whether translators are able to act upon the extent and nature of MT use in post-editing, which is a slightly more constricted approach to this concept than that adopted in previous research (see Läubli and Green in this volume and Olohan 2011 for an analysis of agency in translators' interactions with CAT tools).

A first distinction to make in relation to agency in this context regards the difference between human and automatic post-editing (Simard, Goutte and Isabelle 2007). The core purpose of automatic post-editing is to improve MT outputs automatically in 'black-box' conditions where MT systems cannot be internally adapted or improved. This has usually involved developing post-editing models that can learn from patterns found in aligned parallel data with the raw MT output on one side and the post-edited version on the other (*ibid.*). Once trained, these models can be used to automate the post-editing process to improve the MT output after it has been generated. Although attention has shifted to NMT development, automatic post-editing is nevertheless an active field where new developments are constantly reported (see Astudillo, Graça and Martins 2018). While automatic post-editing is not the focus of this chapter, it can be conceived as the starting point of a human agency spectrum in post-editing tasks. It represents a situation where, from translators' perspective, there is no human involvement in the post-editing process.

As mentioned previously, human post-editing, on the other hand, can involve different degrees of human control. Where MT systems adapt longitudinally to translators' corrections and make different suggestions on the fly based on their edits, translators are not necessarily required to use the MT suggestions to begin with; in some settings, they may simply pick useful suggestions as and when they appear. This is arguably the opposite end of the agency spectrum relative to automatic

post-editing. First, this is because in this context translators are given almost full control over the extent of MT use. Second, where the MT suggestions adapt to translators' edits, this also gives translators more control in shaping the behaviour of the MT system. Figure 1 illustrates this by classing automatic post-editing as 'MT-centred' and interactive and adaptive post-editing as 'human-centred'.

Figure 1. Spectrum of human agency in the post-editing process

Considering Figure 1 is focused strictly on post-editing and does not cover MT development or contexts where MT suggestions are not provided at all, the place of automatic post-editing as the starting point of the spectrum is largely uncontroversial. Static post-editing, on the other hand, is placed in an intermediate position, with more room for variation. It is proposed here that static post-editing can move closer towards either end of the spectrum depending on perspective and on the context of MT use. For instance, in contexts akin to the early days of MT technology where post-editors with no linguistic qualifications could be tasked with improving MT outputs as an alternative to professional translation, static post-editing would move closer towards the MT-centred end of the spectrum, where humans could be seen to assist MT systems. This paradigm would also apply to more recent crowdsourcing models where non-professional crowd workers may be involved in post-editing tasks, sometimes monolingually (see Hu *et al.* 2014 and also Jiménez-Crespo in this volume). Conversely, where professional translators edit static MT to improve their productivity while perhaps using it together with other resources such as translation memories and term bases, this would be closer to the human-centred end of the spectrum, where MT is used as a tool. These possibilities co-exist. In some business models, non-professional bilinguals may improve MT outputs in tasks that are based mainly on MT as a non-optional starting point for consideration^{iv} (also see Yamada in this volume). Post-editing can also be a service that strictly requires qualified or experienced professionals who use MT as one of various sources of 'matches' at hand. This last possibility is not dissimilar from contexts covered by the ISO 18587 standard, which recognizes that post-editing can be carried out in integrated environments including translation memories and other resources. Even in this case, however, translators are recommended to use as much of the MT output as possible, so their level of freedom in deciding how much they should draw on the MT suggestions could still be considered lower than that represented by the 'human-centred' end of the spectrum in Figure 1. A third possibility consists of situations where translators decide to use MT as a productivity-booster or where they are given this option by their employers. Here post-editing is not necessarily classed as a separate service, but as an activity carried out as part of the overall translation process. In this case, static post-editing is even closer to the 'human-centred' paradigm shown in Figure 1. It should be noted that these scenarios are inevitably approximations, since translators' agency in post-editing does not depend only on the nature of the task, but also – importantly – on the client, the nature of the commission, the translation company (if applicable),

among other elements. It is also worth noting that translators' agency is multidimensional and can in some situations simply be invoked in different ways rather to a lesser or higher degree. It may be argued, for instance, that restriction of choice can enhance translators' ability to be inventive in how they address linguistic issues in reaching a compromise between maximum MT use and expectations of high target-text quality. The formulation in Figure 1 nevertheless shows that post-editing is multifaceted.

Indeed, another bone of contention in Figure 1 regards whether the human-centred end of the spectrum should be called 'post-editing' at all. This is regarded here as mostly a matter of nomenclature which is secondary to the potentially higher levels of human control brought about by the interactive paradigm. While conventions in this respect are currently in a state of flux, like the stance adopted by previous researchers (e.g. O'Brien 2016, Bundgaard and Christensen 2019) this chapter considers post-editing and its different contemporary facets to be part of CAT.

Conclusion and Future Directions

The aim of this chapter was to discuss how post-editing of MT has evolved as a practice and service. The chapter regards a move from an MT-centred task to a human-centred one as a key underlying pattern of post-editing's evolutionary process. The way in which post-editing was conceptualised in the early days of MT coupled with the then even more limited capabilities of MT technology has likely helped to nurture negative perceptions of post-editing among professional translators. However, different ways of making use of MT in the human translation process have started to emerge, which has recently diversified perceptions and definitions of post-editing tasks. The chapter responds to these developments by regarding post-editing as an encompassing term that may be used to describe not only static editing of MT outputs but also the use of MT as an additional source of matches in CAT environments. It is also understood here that the use of MT in this way may be offered as a separate service under the auspices of the ISO 18587, or simply as a way of improving translators' productivity in regular commissions where MT use is allowed. Increasingly blurred lines between post-editing and from-scratch translation are suggested as the direction professional translation is most likely to take. Bringing post-editing and from-scratch translation closer together is likely to move translation commissions closer towards the 'human-centred' end of the spectrum shown in Figure 1, which the present author sees as a welcome development.

Given the changing paradigms discussed above, post-editing guidelines focusing on strict levels of editing will probably change in the future as MT technology continues to improve. Even though the chapter discusses research showing that light post-editing is still required to improve MT's usability, the use of raw MT for 'gisting' already is a reality in many settings. Translators' role in light post-editing may therefore become more focused on terminological checks and content sign-off in the future. This does not mean, however, that light post-editing should be abandoned or indeed that it should be reserved for non-professionals. On the contrary, while there may be business models that accommodate non-expert bilinguals and also monolinguals, in the short to medium term hierarchical perceptions of post-editing – and light post-editing in particular – as a lesser service risk creating problematic imbalances in the translation market. The chapter has discussed how post-editing of NMT may pose issues of a slightly different nature that make professional translation expertise even more of a requirement in contexts where information accuracy is an important factor. This too arguably helps to move post-editing tasks towards a human-centred paradigm. This also means that conceptualising tasks, services and modes of human-computer interaction (see Läubli and Green in this volume) in encompassing ways is likely to be more productive than fragmenting and reducing

the range of tasks under translators' aegis. In the same way that technologies are more integrated so too – the present author would argue – should translation services. In practical terms, this means, among other things, that translators would ideally be able to offer a range of services that challenges reductive views of translation as the mere replacement of symbols. This also means that as commissions move closer to the human-centred end of the spectrum in Figure 1, it may make more sense for the terms 'post-editor' and 'translator' to be merged or used interchangeably depending on the job. In other words, as the need for interacting with MT moves from the peripheries to the centre of the translation industry (see Koponen 2016), post-editing is no longer a separate task undertaken only in MT development circles. Rather, it is increasingly part of what *translators* do.

As for research, the use of post-editing as a method for translating literary text is clearly emerging as a trend. This has made translation technology investigations venture into a territory that for years has been off-limits for research in this area. Initial investigations of this subject have presented promising results. These exciting developments may nevertheless require redefining important concepts and methodologies. Going beyond the sentence as a default level of analysis – probably a welcome move in any area – may prove to be paramount in studies on MT post-editing for literary content. Similarly, quality assessment will also need to revisit theoretical paradigms that might not be consistent with the pragmatic approach often taken to quality in MT research involving, for example, automatic evaluations. In addition, more technological integration and interaction between translators and MT systems call for more research on the implications of different working methods. The human-machine interface is still relatively poorly understood in relation to issues such as decision-making, agency and cognitive processing, so research on these subjects should hopefully continue to evolve.

References

- Allen, J. (2003) 'Post-editing', in H. L. Somers (ed.) *Computers and Translation: a translator's guide*, Amsterdam & Philadelphia: John Benjamins Publishing Company, 297-317.
- Alves, F., A. Koglin, B. Mesa-Lao, M. G. Martínez, N. B. d. L. Fonseca, A. d. M. Sá, J. L. Gonçalves, K. S. Szpak, K. Sekino and M. Aquino (2016) 'Analysing the impact of Interactive machine translation on post-editing effort', in M. Carl, S. Bangalore and M. Shaeffer (eds) *New directions in empirical translation process research: exploring the CRITT TPR-DB*, Heidelberg: Springer, 77-94.
- Astudillo, R., J. Graça and A. Martins (eds) (2018) *The 13th Conference of The Association for Machine Translation in the Americas. Workshop Proceedings: Translation Quality Estimation and Automatic Post-Editing*. Available online: <http://aclweb.org/anthology/W18-2100> [last access 12 December 2018].
- Aziz, W., M. Koponen and L. Specia (2014) 'Sub-Sentence level analysis of machine translation post-editing effort', in S. O'Brien, L. W. Balling, M. Carl, M. Simard and L. Specia (eds) *Post-editing of machine translation: processes and applications*, Newcastle upon Tyne: Cambridge Scholars Publishing, 170-199.
- Bar-Hillel, Y. (1951) 'The present state of research on mechanical translation', *American Documentation* 2(4): 229-237.

- Bernth, A. and C. Gdaniec (2002) 'MTranslatability', *Machine Translation* 16(3): 175-218.
- Besacier, L. (2014) 'Traduction automatisée d'une oeuvre littéraire: une étude pilote' [Automatic translation of a literary work: A pilot study], *21ème Traitement Automatique du Langage Naturel (TALN 2014)*. Marseille: 389-394. Available online: <https://hal.inria.fr/hal-01003944/document> [last access 12 December 2018].
- Bundgaard, K. and T. P. Christensen (2019) 'Is the concordance feature the new black? A workplace study of translators' interaction with translation resources while post-editing TM and MT matches', *The Journal of Specialised Translation* 31: 14-37. Forthcoming
- Cadwell, P., S. O'Brien and C. S. C. Teixeira (2017) 'Resistance and accommodation: factors for the (non-) adoption of machine translation among professional translators', *Perspectives*: 1-21.
- Carl, M., B. Dragsted, J. Elming, D. Hardt and A. L. Jakobsen (2011) 'The Process of Post-Editing: a Pilot Study', in B. Sharp, M. Zock, M. Carl and A. L. Jakobsen (eds) *Human-Machine Interaction in Translation*, Copenhagen: Samfundslitteratur, 131-142.
- Castilho, S., F. Alves, S. O'Brien and M. O'Brien (2014) 'Does post-editing increase usability? A study with Brazilian Portuguese as Target Language', *Proceedings of the seventh annual conference of the European Association for Machine Translation*, European Association for Machine Translation, 183-190. Available online: <http://doras.dcu.ie/19997/> [last access 12 December 2018].
- Castilho, S., J. Moorkens, F. Gaspari, I. Calixto, J. Tinsley and A. Way (2017) 'Is neural machine translation the new state of the art?', *The Prague Bulletin of Mathematical Linguistics* 108: 109-120.
- Green, S., J. Chuang, J. Heer and C. D. Manning (2014) 'Predictive Translation Memory: A mixed-initiative system for human language translation', *Proceedings of the 27th annual ACM symposium on User interface software and technology*. Honolulu, HI USA: Association for Computing Machinery, 177-187.
- Green, S., J. Heer and C. D. Manning (2013) 'The Efficacy of Human Post-Editing for Language Translation', in *Proceedings of ACM Human Factors in Computing Systems (CHI)*. Paris: Association for Computing Machinery, 439-448.
- Guerberof, A. (2013) 'What do professional translators think about post-editing?', *The Journal of Specialised Translation* 19: 75-95.
- Hassan, H., A. Aue, C. Chen, V. Chowdhary, J. Clark, C. Federmann, X. Huang, M. Junczys-Dowmunt, W. Lewis, M. Li, S. Liu, T.-Y. Liu, R. Luo, A. Menezes, T. Qin, F. Seide, X. Tan, F. Tian, L. Wu, S. Wu, Y. Xia, D. Zhang, Z. Zhang and M. Zhou (2018) 'Achieving Human Parity on Automatic Chinese to English News Translation', <https://arxiv.org/abs/1803.05567> [last access 12 December 2018].
- Hu, C., P. Resnik and B. B. Bederson (2014) 'Crowdsourced Monolingual Translation', *ACM Transactions on Computer-Human Interaction* 21(4): 22:21-22:35.
- ISO (2017) *ISO 18587 Translation Services – Post-editing of Machine Translation Output – Requirements*, Geneva: International Standards Office.

- Jia, Y., M. Carl and X. Wang (2019) 'How does the post-editing of Neural Machine Translation compare with from-scratch translation? A product and process study', *The Journal of Specialised Translation* 31. Forthcoming
- Kinnunen, T. and K. Koskinen, eds. (2010) *Translators' agency*, Tampere, Tampere University Press.
- Koehn, P., V. Alabau, M. Carl, F. Casacuberta, M. García-Martínez, J. González-Rubio, F. Keller, D. Ortiz-Martínez, G. Sanchis-Trilles and U. Germann (2015) *CASMACAT: Final public report*. Available online: <http://www.casmacat.eu/uploads/Deliverables/final-public-report.pdf> [last access 12 December 2018].
- Koehn, P. and Haddow, B. (2009) 'Interactive assistance to human translators using statistical machine translation methods', in *Proceedings of the 12th Machine Translation Summit (MT Summit XII)*. Ottawa: Association for Machine Translation in the Americas. Available online: <http://www.mt-archive.info/05/MTS-2009-Koehn-2.pdf> [last access 17 December 2018].
- Koponen, M. (2012) 'Comparing human perceptions of post-editing effort with post-editing operations', in *Proceedings of the Seventh Workshop on Statistical Machine Translation*. Montréal: Association for Computational Linguistics, 181-190. Available online: <http://www.statmt.org/wmt12/pdf/WMT23.pdf> [last access 12 December 2018].
- Koponen, M. (2016) 'Is machine translation post-editing worth the effort? A survey of research into post-editing and effort', *The Journal of Specialised Translation* 25: 131-148.
- Krings, H. P. (2001) *Repairing texts: empirical investigations of machine translation post-editing processes*, Kent: Kent State University Press.
- Langlais, P. and G. J. M. T. Lapalme (2002) 'Trans Type: Development-Evaluation Cycles to Boost Translator's Productivity', *Machine Translation* 17(2): 77-98.
- Läubli, S., Sennrich, R. and M. Volk (2018) 'Has Machine Translation Achieved Human Parity? A Case for Document-level Evaluation', in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels: Association for Computational Linguistics, 4791–4796. Available online: <http://aclweb.org/anthology/D18-1512> [last access 18 December 2018].
- Läubli, S. and D. Orrego-Carmona (2017) 'When Google Translate is Better than Some Human Colleagues, those People Are no Longer Colleagues', in *Proceedings of the 39th Conference Translating and the Computer*. London: Aslib, 59-69. Available online: https://www.zora.uzh.ch/id/eprint/147260/1/AS5678529422008321512398099109_content_1.pdf [last access 12 December 2018].
- Lommel, A. (2018) 'Augmented translation: A new approach to combining human and machine capabilities', in *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 2: User Papers)*. Available Online: <http://aclweb.org/anthology/W18-1905> [last access 12 December 2018].
- Massardo, I., J. van der Meer, S. O'Brien, F. Hollowood, N. Aranberri and K. Drescher (2016) *MT Post-Editing Guidelines*, Amsterdam: Translation Automation User society.

- Mitchell, L., J. Roturier and S. O'Brien (2013) 'Community-based post-editing of machine-translated content: monolingual vs. bilingual', in S. O'Brien, M. Simard and L. Specia (eds) *Proceedings of the MT Summit XIV Workshop on Post-editing Technology and Practice*. Available online: <http://doras.dcu.ie/20030/> [last access 12 December 2018].
- Moorkens, J. and S. O'Brien (2015) 'Post-editing evaluations: Trade-offs between novice and professional participants', in Durgar El-Kahlout, M. Özkan, F. Sánchez-Martínez et al. (eds) *Proceedings of European Association for Machine Translation (EAMT)*. Available online: <http://www.aclweb.org/anthology/W15-4910> [last access 12 December 2018].
- Moorkens, J., A. Toral, S. Castilho and A. Way (2018) 'Translators' perceptions of literary post-editing using statistical and neural machine translation', *Translation Spaces* 7(2): 240–262.
- Nitzke, J. (2016) 'Monolingual post-editing: An exploratory study on research behaviour and target text quality', in S. Hansen-Schirra and S. Grucza (eds) *Eyetracking and Applied Linguistics*, Berlin: Language Science Press, 83-109.
- O'Brien, S. (2016) 'Post-Editing and CAT', *EST Newsletter* 48.
- O'Brien, S. (2011) 'Towards predicting post-editing productivity', *Machine Translation* 25(3): 197-215.
- Olohan, M. (2011) 'Translators and translation technology: The dance of agency', *Translation Studies* 4(3): 342-357.
- Papineni, K., S. Roukos, T. Ward and W. J. Zhu (2002) 'BLEU: a method for automatic evaluation of machine translation', in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Philadelphia: Association for Computational Linguistics, 311-318. Available online: <https://www.aclweb.org/anthology/P02-1040.pdf> [last access 12 December 2018].
- Plitt, M. and F. Masselot (2010) 'A productivity test of statistical machine translation post-editing in a typical localization context', *Prague Bulletin of Mathematical Linguistics* 93: 7-16.
- Reifler, E. (1952) 'Mechanical translation with a pre-editor and writing for mechanical translation', presented at the Conference on Mechanical Translation, 17-20 June, Cambridge. Available online: <http://www.mt-archive.info/MIT-1952-Reifler-1.pdf> [last access 12 December 2018].
- Schwartz, L. (2014) 'Monolingual post-editing by a domain expert is highly effective for translation triage', in S. O'Brien, M. Simard and L. Specia (eds) *Proceedings of the Third Workshop on Post-Editing Technology and Practice*, Vancouver: Association for Machine Translation in the Americas, 34-44.
- Screen, B. (2019) 'What effect does post-editing have on the translation product from an end-user's perspective?', *The Journal of Specialised Translation* 31. Forthcoming
- Simard, M., C. Goutte and P. Isabelle (2007) 'Statistical Phrase-based Post-editing', in *Proceedings of NAACL HLT 2007*, Madison: Association for Computational Linguistics/Omnipress Inc, 508-515.

- Specia, L. (2011) 'Exploiting objective annotations for measuring translation post-editing effort', in M. L. Forcada, I. Depraetere and V. Vandeghinste (eds) *Proceedings of the 15th International Conference of the European Association for Machine Translation*, Leuven: European Association for Machine Translation, 73-80. Available online: <http://www.mt-archive.info/EAMT-2011-Specia.pdf> [last access 12 December 2018].
- Specia, L., F. Blain, V. Logacheva, R. Astudillo and A. Martins (2018) 'Findings of the WMT 2018 Shared Task on Quality Estimation', in *Proceedings of the Third Conference on Machine Translation (WMT)*, Brussels: Association for Computational Linguistics, 689-709. Available online: <http://aclweb.org/anthology/W18-6451> [last access 12 December 2018].
- Tatsumi, M. and J. Roturier (2010) 'Source text characteristics and technical and temporal post-editing effort: what is their relationship?', in V. Zhechev (ed.) *Proceedings of the Second Joint EM+/CNGL Workshop 'Bringing MT to the User: Research on Integrating MT in the Translation Industry'*, Denver, 43-51. Available online: <https://pdfs.semanticscholar.org/5573/e1c9a9466c2b84074f1c9a1f3e4387570c40.pdf> [last access 12 December 2018].
- Toral, A., S. Castilho, K. Hu and A. Way (2018) 'Attaining the Unattainable? Reassessing Claims of Human Parity in Neural Machine Translation', <https://arxiv.org/abs/1808.10432>.
- Toral, A. and A. Way (2018) 'What Level of Quality can Neural Machine Translation Attain on Literary Text?', in J. Moorkens, S. Castilho, F. Gaspari and S. Doherty (eds) *Translation Quality Assessment: from Principles to Practice*. Cham: Springer, 263-287.
- Underwood, N., B. Mesa-Lao, M. G. Martínez, M. Carl, V. Alabau, J. González-Rubio, L. A. Leiva, G. Sanchis-Trilles, D. Ortiz-Martínez and F. Casacuberta (2014) 'Evaluating the Effects of Interactivity in a Post-Editing Workbench', in N. Calzolari, K. Choukri, T. Declerck et al. (eds) *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik: European Language Resources Association, 553-559.
- van Egdom, G.-W. and M. Pluymaekers (2019) 'Why go the extra mile? How different degrees of post-editing affect perceptions of texts, senders and products among end users', *The Journal of Specialised Translation* 31. Forthcoming.
- Vermeer, H. J. (2000) 'Skopos and commission in translational action'.in L. Venuti (ed.) *The Translation Studies Reader*, London: Routledge, 221-232.
- Vieira, L. N. (2014) 'Indices of Cognitive Effort in Machine Translation Post-Editing', *Machine Translation* 28(3-4): 187-216.
- Vieira, L. N. (2017) 'From process to product: links between post-editing effort and post-edited quality', in A. L. Jakobsen and B. Mesa-Lao (eds) *Translation in Transition: Between Cognition, Computing and Technology*, Amsterdam: John Benjamins Publishing Company, 162-186.
- Vieira, L. N. and E. Alonso (2018) *The use of machine translation in human translation workflows: Practices, perceptions and knowledge exchange*, Milton Keynes: Institute of Translation and Interpreting. Available online: <https://www iti.org.uk/images/downloads/ITIRReport-Lucas.pdf> [last access 12 December 2018].

Wagner, E. (1985) 'Post-editing Systran - A challenge for Commission Translators', *Terminologie et Traduction* 3: 1-7.

Yamada, M. (2019) 'The impact of Google Neural Machine Translation on Post-editing by student translators', *The Journal of Specialised Translation* 31. Forthcoming.

Zhechev, V. (2014) 'Analysing the post-editing of Machine Translation at Autodesk', in S. O'Brien, L. W. Balling, M. Carl, M. Simard and L. Specia (eds) *Post-editing of Machine Translation: Processes and Applications*, Newcastle: Cambridge Scholars Publishing, 2-24.

Further reading

Guerberof, A. and J. Moorkens (2019) 'Machine translation and post-editing training as part of a master's programme', *The Journal of Specialised Translation* 31. Forthcoming.

This article deals with important aspects of incorporating post-editing into translator training, an area of research that is not covered by the present chapter.

Carl, M., S. Bangalore and M. Shaeffer (eds) (2016) *New directions in empirical translation process research: exploring the CRITT TPR-DB*, Heidelberg: Springer.

Readers are directed to several chapters in this edited volume covering aspects of post-editing that could not be discussed in detail above. In particular, chapter 6, by Daems et al., analyses how reference materials are used in from-scratch translation and post-editing. Chapters 11 and 14, by Schmaltz et al. and Nitzke and Oster, respectively, provide useful methodological insights into how the CRITT (Centre for Research and Innovation in Translation and Translation Technology) translation process research database can be used in post-editing research.

O'Brien, S., L. W. Balling, M. Carl, M. Simard and L. Specia (eds) (2014) *Post-editing of machine translation: processes and applications*, Newcastle upon Tyne: Cambridge Scholars Publishing.

This edited volume covers a range of topics from industrial and academic perspectives. Attention is drawn to chapter 5, by Teixeira, on how translation metadata is used in translation tools.

ⁱ Examples of other studies examining the feasibility of post-editing relative to 'from-scratch' translation include Carl *et al.* (2011) and Jia, Carl, and Wang (2019). For a summary, see Screen (2019).

ⁱⁱ See <https://lilt.com/>.

ⁱⁱⁱ See <https://lilt.com/kb/memory/mt>

^{iv} See <https://unbabel.com/translators/> for a similar example.