

Chapter 5

How to deal with errors in machine translation: Post-editing

Sharon O'Brien

Dublin City University

Machine Translation output can be incorrect, containing errors that need to be fixed, especially if the text is destined for publication and if it is important that it contains no errors. The task of identifying and fixing these errors is called *post-editing* (PE). In this chapter, I provide an overview of the PE process, drawing on both academic and industry sources. I explain how PE is generally divided into *light* and *full* PE, and describe standard guidelines for each type, homing in on issues that arise in the application of this classification. The chapter also surveys the various types of interface used in PE (including word processing and spreadsheet software, and professional computer-aided translation tools), and modes of interaction (traditional, adaptive or interactive). Finally, concepts and tools used by researchers into PE are described, and particular focus is put on the measurement of temporal, technical and cognitive effort.

1 Definition

Machine Translation (MT) is an imperfect technology. For one sentence it might produce an accurate and contextually acceptable translation, but the next sentence might have a serious error in meaning, an omission, an addition, or a stylistic problem. If MT is being used just to obtain the gist of the meaning from a text, there may be no need to fix such errors. However, if MT is being used to create a text for publication or widespread circulation within or outside an organisation, it is usually necessary to fix any errors in the text. The identification of such errors and their revision, or correction, is known as *post-editing*. The term was



already used in the early stages of development of MT systems, when technology was somewhat slower, or certainly less instantaneous than it is now. A text would be sent electronically to the MT system, the system would translate, the text would be returned to the sender who would then edit it, post – or after – the automatic translation stage. The term has stuck, and is even used to describe contemporary processes in which error fixing happens at the same time as automatic translation (see §3.1).

Post-editing is a bilingual language processing task. It is typically, though not exclusively, undertaken by experienced professional translators. When a person is engaged in this task, they are usually referred to as a *post-editor*. However, the usefulness of this name is debatable, especially because the professional translator's typical work environment involves the use of a computer-aided translation tool that combines technologies such as translation memory, terminology management and machine translation (see Kenny 2022: §4 [this volume] and §3 below). Within this working environment, the translator might be editing a fuzzy match for one sentence, translating the next sentence, and post-editing the one that follows. Does the translator change from being a revisor, to a translator, to a post-editor for each of these sentences? Not really! Essentially, the task is still all about translation and revision. The main difference lies in the technological input or support the translator is using from one moment to the next.

When post-editing, the translator has to understand the source language sentence and the target language proposal from the MT system. They have to then identify any errors, devise a strategy for fixing them, and implement those revisions. Fundamentally, post-editing is a revision task.

MT systems can produce a variety of error types, ranging from grammatical errors, to syntactic errors, to unnecessary additions or omissions, to errors in lexical or terminological choice, errors in collocation or style. The type and number of errors that might exist in a text produced by an MT system will vary depending on many factors, such as the language pair and direction, the content type, and the data or techniques used to train the MT engine.

Here are just a few examples in English of some of these error types and their fixes.

(1) Grammatical error

- a. The cat is very protective of her kittens. She scratches anyone *which* tries to touch them.
- b. The cat is very protective of her kittens. She scratches anyone *who* tries to touch them.

- (2) Lexical error
 - a. The cat is very protective of her *pups*. She scratches anyone who tries to touch them.
 - b. The cat is very protective of her *kittens*. She scratches anyone who tries to touch them.
- (3) Syntactic error
 - a. The new-born cygnets on the lake *swam*.
 - b. The new-born cygnets *swam* on the lake.
- (4) Collocation error
 - a. The house had no *flowing water*.
 - b. The house had no *running water*.

2 Levels of post-editing and guidelines

One of the main objectives of MT is to enable more information to be translated between more languages, but to achieve this more rapidly than is possible through translation unaided by MT and, indeed, at a lower cost. These objectives have led to a distinction between different levels of post-editing, mainly characterised as *light post-editing* or *full post-editing*.

Light post-editing is understood to mean that only essential fixes should be implemented and that this should be done rapidly. Full post-editing, on the other hand, means that all errors in the MT output should be fixed, and this is expected to take more time when compared with “light” post-editing. For both levels, it is generally expected that the translation can be produced faster than translation without computer-assisted translation tools. The International Standards Organisation (ISO) has produced a standard for post-editing known as “ISO 18857:2017” (ISO 2017). In this standard light post-editing is described as the “process of post-editing to obtain a merely comprehensible text without any attempt to produce a product comparable to a product obtained by human translation” (p.2). Full post-editing is defined as the “process of post-editing to obtain a product comparable to a product obtained by human translation” (ibid.).

These definitions are, however, conceptually problematic. In the first instance, it is difficult to articulate what exactly the differences are between the two, except in the very general way we have described them above. What, exactly, is an “essential” fix? This might differ from one organisation to the next, depending on requirements. What does “merely comprehensible” mean and how would that

be measured? And can light post-editing really be done without *any attempt* to produce something comparable to human translation? Furthermore, it is unclear how much longer full post-editing would take compared with light post-editing. Such questions prompted the translation industry to describe these levels more formally. For example, the Translation Automation User Society (TAUS) created guidelines suggesting that full post-editing would include stylistic changes, whereas light post-editing would not.

The TAUS guidelines (TAUS 2010) for light and full post-editing are listed in Table 1, where comparable guidelines appear side by side, and an empty cell indicates that there is no comparable guideline in one set of guidelines:

According to the ISO18857 standard (p.6), the objectives of post-editing are to ensure:

- Comprehensibility of the post-edited output;
- Correspondence of source language content and target language content;
- Compliance with post-editing requirements and specifications defined by the TSP.

where “TSP” stands for “translation service provider” and is defined as a “language service provider that delivers translation services” (p.4).

These objectives can be attained by ensuring that the following criteria are met (p.6):

- Terminology/lexical consistency, as well as compliance with domain terminology;
- Use of standard syntax, spelling, punctuation, diacritics, special symbols and abbreviations and other orthographical conventions of the target language;
- Compliance with any applicable standards;
- Correct formatting;
- Suitability for the target audience and for the purpose of the target language content;
- Compliance with client-TSP agreements.

Table 1: The TAUS post-editing guidelines.

Light post-editing	Full post-editing
Aim for semantically correct translation.	Aim for grammatically, syntactically and semantically correct translation.
	Ensure that key terminology is correctly translated and that untranslated terms belong to the client's list of "Do Not Translate" terms.
Ensure that no information has been accidentally added or omitted.	Ensure that no information has been accidentally added or omitted.
Edit any offensive, inappropriate or culturally unacceptable content.	Edit any offensive, inappropriate or culturally unacceptable content.
Use as much of the raw MT output as possible.	Use as much of the raw MT output as possible.
Basic rules regarding spelling apply.	Apply basic rules regarding spelling, punctuation and hyphenation.
No need to implement corrections that are of a stylistic nature only.	
No need to restructure sentences solely to improve the natural flow of the text.	
	Ensure that formatting is correct.

There are overlaps between the two sets of guidelines, though they prioritise different aspects of the task. Taken together, they represent typical guidelines for post-editing. The TAUS guidelines encourage as much re-use of raw MT output as is practical, whereas the ISO guidelines focus more on agreements, standards and suitability for the target audience. The notion of reusing as much of the raw MT output as is possible is an essential aspect of the post-editing task. It is very easy for a translator to simply ignore the MT output, delete it and translate the source sentence directly. In fact, many translators are tempted to do this because

they believe that they can produce a better translation and that it will take less time than post-editing. While the first belief was certainly true some time ago, the development of neural machine translation has, in general, increased the quality of MT to such an extent that raw output is now much more useful and usable. The idea that translation would take less time than post-editing is, on the other hand, open to debate. Studies have shown that post-editing can certainly be faster than translation, even if translators think that they are faster (e.g. Guerbero of Arenas 2014). Being able to rapidly assess the output from an MT system, decide if it is usable, and what edits are required is something that can be honed with practice.

Levels of post-editing are conceptually linked with levels of quality, though as will be shown in the examples below this linkage is problematic. Light post-editing is seen to be linked with “good enough quality”, or text that is “merely comprehensible”, i.e. text that should be accurately translated, but that does not necessarily have to flow very naturally or be stylistically sophisticated. On the other hand, full post-editing is linked with “quality similar or equal to human translation”. Here again, we run into some difficulty because the inherent assumption is that “human translation” is *always* of a high standard, something that is frankly not always the case.

To better understand the complex, and sometimes confusing, relationship between levels of post-editing and levels of quality, let us look at an example:

- (5) a. In a new report on the quality of teaching practice, inspectors said pre-school teachers could be *training* to integrate languages such as French, German and Polish in early *learn* settings.
- b. In a new report on the quality of teaching practice, inspectors said pre-school teachers could be *trained* to integrate languages such as French, German and Polish in early *learning* settings.

The two errors in (a) are quickly fixed to render sentence (b). On the one hand, we could say that we have edited (a) *lightly*; we implemented two rapid edits. However, if we were to follow the light post-editing guidelines, we probably would not implement any edits at all. Implementation of “full post-editing” guidelines would mean that the two errors must be fixed to produce a translation that is semantically, syntactically and grammatically correct. So, with these two rapid edits, have we engaged in light or full post-editing?

Let us take a look at another example.

- (6) a. In addition, it *reiterates the instinct* of modern *language* into the primary school curriculum *as a fitting step*.

- b. In addition, it says *re-instating modern languages* to the primary school curriculum *would be a timely move*.¹

Sentence (a) is rather unclear and so significant editing is required to make it at least semantically correct. Fundamentally, the level of editing depends on the starting point (the quality of each sentence proposed by the MT system) and the targeted quality. The post-editor will pivot between quick edits and more significant edits depending on each sentence and the ultimate quality objective. What is most important here is that the final text has a coherent quality, regardless of what the starting point is for each sentence, and that the quality meets with the translation commissioner's and end users' expectations.

One final issue with the concepts of light and full post-editing is that professional translators generally do not want to agree to produce "merely comprehensible" quality, and no commissioner of translation is really willing to admit that they opted *only* for a light edit.

It is important to know about these concepts and the guidelines that come with them. However, as is hopefully evident, they are not without their challenges.

3 Post-editing interfaces

At a basic level, post-editing can be done in any text editor where the source text is visible and the "raw" MT output can be revised. This could even be a spreadsheet, with the source text in one column and the MT output in an adjacent one. These days, however, professional translation is typically done using computer-aided translation (CAT) environments, especially translation memory (TM) tools. As indicated by Kenny (2022: §4 [this volume]), translation memory is a database that stores segments of texts that have been previously translated. A TM tool is the software application that is used to access, edit and update the text in this database. MT, as is evident from the other chapters in this book, is a different type of technology, though the two are inevitably linked because contemporary data-driven MT systems typically use the data stored in TMs as an important input for machine learning. Additionally, seeing as TM tools are so commonly used by translators in their daily work, MT technology is now linked to, if not completely embedded in, TM tools such as Trados Studio, MQM, MateCat, to name but a few. From a practical perspective, this means that post-editing is frequently carried out in a TM editing environment.

¹<https://www.irishtimes.com/news/education/foreign-languages-could-be-taught-in-preschool-and-primary-department-1.4270886>

There are different ways in which MT can be embedded in a TM environment. For example, the MT system might be called on by the translator if there is no useful match offered from the TM. Alternatively, the translator could customise their TM tool settings to ensure that an MT suggestion is presented automatically when a specific set of conditions are met (for example, when no match is found in the TM).

It is generally accepted that, if the TM is well-maintained and kept up to date, an exact match from the TM database is more valuable to the translator than one from the MT system because the latter might contain errors whereas the former (ideally) should not. It used to also be assumed in industry circles that a “fuzzy match” from the TM of 75% similarity or higher was better than an MT suggestion, for the same reason. MT has, however, become more advanced recently and the suggestions may be very useful to a translator, possibly even more so than a 75% match from the TM. Thus, depending on the TM, the MT system, the language pair and the topic, a translator might customise their TM tool settings so that they see both TM matches and MT suggestions at the same time. In fact, they could even opt to have suggestions from multiple MT systems presented simultaneously. The translator can select the suggestion that they deem to be most useful and edit it as required.

The benefit of using TM/MT integration is that the translator gets more choice and assistance with the translation process. There are some challenges too. With this set up, a lot of information is presented to the translator, which means that they have to process all of this information (the TM match and its match value, the MT suggestion(s), possibly also terminology, and meta data such as who created the TM match and when etc. (see Teixeira & O'Brien 2017), while making a decision on which match to work with and what needs to be revised from a linguistic perspective. This can lead to a high cognitive demand and even overload, and possibly explains why some translators report that post-editing is more demanding and tiring than other forms of translation and revision.

A further challenge with this mixed interface is presented by the post-editing guidelines mentioned above. The translator needs to keep those guidelines in mind when they are working with an MT suggestion, but if they are working with a match from the TM, then they are dealing with a translation generated by a translator (as opposed to an MT system) and so are no longer officially “post-editing”. This mix of post-editing, revising and translating potentially makes the task quite complex and we should also acknowledge that the task is typically done under pressure of time too!

3.1 Traditional, adaptive and interactive post-editing

Above, we have outlined the typical interface used for post-editing in a traditional translation setting using a TM tool, with the MT match appearing as a static suggestion that can be selected and edited or discarded in favour of a TM match or full translation of the source sentence. As with all technologies, new inventions appear and we now have some variations on the default, traditional set up. Two such innovations are adaptive MT and interactive MT.

Adaptive MT is a feature developed by TM/MT integrators whereby the MT system learns *in real time* from the edits implemented by the translator. This tackles one of the weaknesses of MT that translators previously found frustrating, i.e. they would fix an error generated by the MT system, but the system would not learn from their revisions. Therefore, if the same sentence occurred again later on, the post-editor had to fix the error again. The integration with translation memory tackled this problem to some extent: Once a sentence was fixed in a TM environment, it would be saved to the TM database. If the same sentence reoccurred later on, it would be presented to the translator as an exact match from the TM database. However, outside that TM environment with that exact database, the MT system would reproduce the error. With adaptive MT, the MT system “learns” from the edit and, theoretically, should adapt so that the same error is not produced by the MT system again. An example of a tool where this feature has been implemented is Trados Studio.

Interactive MT can be seen as a special form of adaptive MT. Whereas the default way of working with MT is as described above, where the MT system pre-translates the entire segment and it is then presented in full to the translator, interactive MT reacts to each decision the translator makes word by word, phrase by phrase. As the translator accepts or confirms a word, the MT system adjusts the output in real time. This is a considerably different way of interacting with MT output, but it is similar to the concept of predictive texting, to which many people have become accustomed. Lilt is an example of a tool where this is one of the primary features. In fact, Lilt is positioned explicitly as an adaptive *and* interactive interface.² (For other examples of interactive MT, see Torregrosa Rivero 2018.)

Interactive MT calls into question the term “post-editing”. As explained earlier, this term came into use many decades ago when the editing was always done *after* the full text had been translated by the MT system and returned to the requester. With the interactive mode, the machine translation happens in real time, changing as the translator makes decisions in the current moment. As

²<https://lilt.com>, last accessed June 2022

a result, the prefix “post” seems irrelevant. “Interactive MT” is a more accurate term. As happens in many domains, the term “post-editing” is now well established so it may not disappear soon, but it will probably become defunct as time goes by. The task itself – interacting with and fixing MT output – is less likely to become defunct in the near future.

3.2 Research interfaces

Above I described the different mainstream interfaces used for post-editing, as well as the modes of interaction. There are many other interfaces one can use, but the features are essentially the same. As post-editing is a relatively new task for professional translators, there has been a significant amount of interest in it from a research perspective. Research has focused on the types of edits that are typically implemented, on measuring whether the task takes less time than translation, the quality produced, and, not least, on the cognitive processes involved.

To capture data on these topics, researchers have developed their own interfaces for post-editing. The main motivations for doing so are that commercial tools can be too expensive for research projects, their features are relatively sophisticated, sometimes actually acting as a hindrance to the research objectives, or they cannot be controlled enough for experimental research conditions. Examples of two research tools developed for post-editing experiments are Casmacat and Translog. Casmacat was developed as part of an EU-funded research project.³ The aim of this project was to build the next generation translator’s workbench using interactive and adaptive machine translation and to build models of the cognitive processes involved in interactive MT. Translog was initially developed for research into translation processes by the Copenhagen Business School. Translog II is a more recent version of the tool that enables research into translation and post-editing.⁴ Both of these tools can be integrated with technologies that are very useful for research purposes, e.g. key logging, where the keyboard activity is recorded, and eye tracking, which records eye movements and cognitive load during the translation or post-editing process. The use of these tools has helped us better understand post-editing as a task. Nonetheless, while they are really useful for research purposes, they are considerably pared down in terms of features compared with standard commercial TM tools, and so can provide only limited understanding of the task as it is performed in real production environments.

³Co-funded by the European Union under the Seventh Framework Programme Project 287576 (ICT-2011.4.2) - <http://www.casmacat.eu>

⁴<https://sites.google.com/site/centretranslationinnovation/translog-ii?authuser=0>

4 Measuring post-editing effort

When technology advances, it moves from the lab to the public domain, where it is tested for usefulness. Initially, this can act as a disruptor – processes used and accepted for years are disturbed, which, in turn, understandably leads to questions, worry, or irritation. When we are used to doing a task one way, we find it difficult to quickly embrace a new way of doing it. Moreover, if the technology creates new problems for us, we are not likely to embrace it fully. This is the case with MT and post-editing for professional purposes. Consequently, there has been much interest in measuring the effort involved in post-editing to verify, or debunk, the claim that it is “faster” than translating without such automation.

A vast amount of research has been produced in the past fifteen years or so to investigate this question (Koponen 2016). There are too many studies to mention individually here. However, one of the seminal texts on the topic deserves mention because it set the standard for measuring post-editing effort. That text was written in German by Hans-Peter Krings, with the English translation appearing in 2001 (Krings 2001). Krings investigated the effort involved in post-editing compared with translation. He did so at a time when MT produced much lower quality output than it does nowadays and his experimental set-up was necessarily naïve, given the era in which it was conducted: the tasks were done on paper and he used a camera to record the process. What is most important about Krings’ study is not his set up nor his findings, but that he argued for the measurement of effort across *three dimensions* – temporal, technical and cognitive.

All too often, measurement of PE effort focuses only on the temporal dimension, i.e. how long does the task take in comparison with another task. Time is, of course, one of the most important aspects, especially in a commercial environment. Time is relatively easy to measure too, so it tends to be the main focus when commercial organisations wish to measure PE effort.

Krings’ work, however, demonstrates that the other two dimensions also need to be considered. Technical effort measures keyboard and mouse actions, i.e. how many words or parts of words are deleted, added, how many phrases are selected, cut and pasted to another location in the text, etc. Translog, mentioned above, is one tool that enables keyboard logging. As PE is textual revision, understanding the effort involved in the mechanical changes implemented is important. Not only that, but this kind of revision – which involves deletion, re-typing, copying and pasting – requires a lot of keyboard and mouse usage, which is physically tiring and can even lead to strain in the hands and wrists. On the other hand, if the MT output is relatively good, MT can reduce the amount of typing a translator has to do.

Apart from keyboard logging, technical effort is also measured using what is called *edit distance* metrics. Put simply, edit distance counts the minimum number of operations required to transform one string of text into another. An "operation" could be deletion of a word, insertion of a word, or movement of a word or phrase to another location. There are several metrics for measuring edit distance, each of which counts the operations slightly differently. One basic metric is called the *Levenshtein distance*. It counts the minimum number of character insertions, deletions or substitutions needed to change to transform one word, phrase or sentence into another.

For example: Take the word *drink* and the word *drunk*. How many characters have to change to transform one into the other? One: 'i' is substituted by 'u'. Let us make this a bit more complex: if we transform the phrase "He drinks" into "He is drinking", the Levenshtein distance is 6. (Insert 'i', 's' and one space character after 'He'; substitute 'i' for 's' at the end of 'drinks' and insert 'n' and 'g'.)⁵ More sophisticated edit distance measures can be deployed and one that is often used to measure PE edit distance is called *TER*, or the *Translation Edit Rate* (Snover et al. 2006). This can be measured on a scale of 0–1 or 0% to 100%. The lower the score, the lower the PE effort. For example, a score of 30% means, approximately, that 30% of the raw MT output was edited to create the post-edited version of a text string. Challenges exist regarding how best to calculate edit distance and consequently there are several different approaches, with different metrics being proposed on a regular basis.

Temporal and technical effort are relatively easy to measure. Measuring the third dimension – cognitive effort – is much more complex. Cognitive effort refers to hidden cognitive processes such as reading, understanding, comparing source language meaning to that of the MT output, decision making, while taking into account the guidelines and expectations, and monitoring the text as it is revised. These processes take place in the brain and cannot be seen or measured directly. Nonetheless, cognitive effort is still an important aspect to consider. Post-editing is sometimes reported as being more demanding a task than translation without MT as an aid. This is probably due to the list of processes mentioned above and also to the fact that it is a relatively new task for some. Even if translators can produce text faster with MT, they may feel more tired than they would do if they were to produce the translation themselves. Working faster suits commercial production, but not if it results in translator burnout, and that is why cognitive demand is important to consider when measuring PE effort.

But how can we measure cognitive effort? In fact, this is a question for anyone who seeks to measure cognitive effort for any task. Sometimes the effort can be

⁵These calculations can be done online using, for example, <https://planetcalc.com/1721/>.

estimated by asking the person who performs the task to “think aloud” as they work. By doing this, they can highlight cognitive difficulties they encounter. Of course, thinking aloud as you work interferes with and slows down the task itself, so there are disadvantages to this technique. An alternative approach is to record the task on the computer screen as it unfolds, then to replay that as a video when the task is completed, and ask the task performer to retrospectively discuss the problems they encountered. This has the advantage of not slowing down the task itself, but it has the disadvantage that the person may not remember all of the issues they encountered. Finally, researchers have attempted to measure cognitive effort in post-editing using eye tracking, a technology that records where the eyes fall on the screen, as well as how long the eyes rest on parts of the text (called *fixation duration*), and even the *pupil dilation*, a measurement of pupil size. These are known to be good measures of cognitive effort. Yet, the challenges are obvious: you need expensive eye tracking technology, sophisticated knowledge in how to use it and interpret the data it produces, and you need to control the data collection environment so that users do not move their heads too much, or the light does not change substantially because this affects the pupil size, and so on. Since measuring cognitive effort is a considerable challenge, understandably few include it when they measure PE effort. Nonetheless, it is important to recognise cognitive effort as an essential component of the effort involved in post-editing.

There is a final note to add here on measuring PE effort. The amount of effort should indirectly tell us something about the quality of the output produced by a specific MT system, for a language pair and topic. Therefore, we can use PE effort as a form of MT quality evaluation. The lower the quality from the MT system, the more changes and time will be required. MT quality can be measured in other ways, by, for example, identifying, classifying and counting the number of errors produced. This is a useful form of MT quality evaluation but taking the PE effort into consideration is potentially even more informative because it reveals how easy or difficult it is to work with the MT output to produce a defined level of quality.

5 Post-editor profiles and training

What makes a good post-editor? And what kind of training should be provided? As MT has become more of a mainstream technology, these are two questions that have preoccupied those in the language industry as well as in academia (see, for example, Nitzke & Hansen-Schirra 2021).

Taking the first question, a common suggestion is that to be a good post-editor, one first has to be a good translator. Intuitively we know that some people are great translators but not great revisors and vice versa. By deduction then, some will be good post-editors and others not.

But what does it mean to be a good post-editor? This question has received some attention already. de Almeida & O'Brien (2010), for example, suggest that a good post-editor has:

1. The ability to identify issues in the raw MT output that need to be addressed and to fix them appropriately;
2. The ability to carry out the post-editing task with reasonable speed, so as to meet the expectations of daily productivity for this type of activity;
3. The ability to adhere to the guidelines, so as to minimise the number of "preferential" changes, or changes that are not strictly speaking necessary, and which are normally outside the scope of PE.

It could be argued that point (1) means that the post-editor must first have mastered translation skills. Points (2) and (3) suggest that a post-editor needs to be able to work quickly and to adhere to guidelines, resisting the temptation to over-edit. Ultimately, being a "good" post-editor is closely linked with an individual's attitude towards MT as a technology (see Guerberoof Arenas 2013 for a deeper discussion). If a translator dislikes MT as a technology, he or she will possibly be tempted to delete or ignore every MT suggestion. Assuming this in turn leads to more time being required for the translation task, as well as a higher cost for the commissioner, then that person cannot be classified as a "good" post-editor, though, again, much depends on the context and the quality produced by the MT system in the first instance.

As for the second question, there has been a growing focus on training in the past decade. As MT was slowly integrated into other CAT tools and translation production processes, professional translators needed training in the form of continuous professional development, e.g. workshops dedicated to learning about MT and post-editing. Over time, MT and PE have been incorporated into translator training programmes in universities. There are many different approaches to this training. Some universities offer entire stand-alone courses on post-editing, some incorporate it into revision courses, and others weave it into translation technology courses (see O'Brien & Vázquez 2019).

A core focus of this training is ensuring that translation students understand the most recent approaches to MT, its strengths and its limitations, how to evaluate it and how to post-edit. Importantly, understanding when and how MT ought

to be used has become central to training, both for translation students and for those who are not trained in translation (see Bowker & Ciro (2019) for a discussion of “MT literacy”).

References

- Bowker, Lynne & Jairo Buitrago Ciro. 2019. *Machine translation and global research*. Bingley: Emerald Publishing.
- de Almeida, Gisele & Sharon O’Brien. 2010. Analysing post-editing performance. Correlations with years of translation experience. In *Eamt 2010. Proceedings of the 14th annual conference of the European Association for Machine Translation*. EAMT. <http://www.mt-archive.info/10/EAMT-2010-Almeida.pdf>.
- Guerberof Arenas, Ana. 2013. What do professional translators think about post-editing? *Journal of Specialised Translation* 19. 75–95. https://www.jostrans.org/issue19/art_guerberof.php.
- Guerberof Arenas, Ana. 2014. Correlations between productivity and quality when post-editing in a professional context. *Machine Translation* 28. 165–186.
- ISO. 2017. *ISO 18857:2017. Translation services – post-editing of machine translation output: Requirements*. <https://www.iso.org/standard/62970.html>.
- Kenny, Dorothy. 2022. Human and machine translation. In Dorothy Kenny (ed.), *Machine translation for everyone: Empowering users in the age of artificial intelligence*, 23–49. Berlin: Language Science Press. DOI: 10.5281/zenodo.6759976.
- Koponen, Maarit. 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. https://www.jostrans.org/issue25/art_koponen.pdf.
- Krings, Hans P. 2001. *Repairing texts: empirical investigations of machine translation post-editing processes*. Ohio: Kent State University Press.
- Nitzke, Jean & Silvia Hansen-Schirra. 2021. *A short guide to post-editing*. (Translation and Multilingual Natural Language Processing 16). Berlin: Language Science Press. DOI: 10.5281/zenodo.5646896.
- O’Brien, Sharon & Silvia Rodríguez Vázquez. 2019. Translation and technology. In Sara Laviosa & Maria González-Davies (eds.), *Routledge handbook of translation and education*, 264–277. London: Routledge.
- Snover, Matthew, Bonnie Dorr, Rich Schwartz, Linnea Micciulla & John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th conference of the Association for Machine Translation in the Americas: Technical papers*, 223–231. Cambridge, Massachusetts: Association for Machine Translation in the Americas. <https://aclanthology.org/2006.amta-papers.25/>.

- TAUS. 2010. *MT post-editing guidelines*. <https://www.taus.net/academy/best-practices/postedit-best-practices/machine-translation-post-editing-guidelines>.
- Teixeira, Carlos S. C. & Sharon O'Brien. 2017. Investigating the cognitive ergonomic aspects of translation tools in a workplace setting. *Translation Spaces* 6(1). 79–103.
- Torregrosa Rivero, Daniel. 2018. *Black-box interactive translation prediction*. PhD Thesis. Alicante: University of Alicante. https://rua.ua.es/dspace/bitstream/10045/77110/1/thesis_daniel_torregrosa_rivero.pdf.