The Translator's Extended Mind

Yuri Balashov

Minds and Machines
Journal for Artificial Intelligence,
Philosophy and Cognitive Science

ISSN 0924-6495
Volume 30
Number 3

Minds & Machines (2020) 30:349-383
DOI 10.1007/s11023-020-09536-5

Springer

Minds and Machines (2020) 30:349–383
https://doi.org/10.1007/s11023-020-09536-5

GENERAL ARTICLE

The Translator's Extended Mind

Yuri Balashov[1]

Received: 31 May 2020 / Accepted: 14 August 2020 / Published online: 19 September 2020

Abstract

The rapid development of natural language processing in the last three decades has drastically changed the way professional translators do their work. Nowadays most of them use *computer-assisted translation* (CAT) or *translation memory* (TM) tools whose evolution has been overshadowed by the much more sensational development of *machine translation* (MT) systems, with which TM tools are sometimes confused. These two language technologies now interact in mutually enhancing ways, and their increasing role in human translation has become a subject of behavioral studies. Philosophers and linguists, however, have been slow in coming to grips with these important developments. The present paper seeks to fill in this lacuna. I focus on the semantic aspects of the highly distributed human–computer interaction in the CAT process which presents an interesting case of an *extended cognitive system* involving a human translator, a TM tool, an MT engine, and sometimes other human translators or editors. Considered as a whole, such a system is engaged in representing the linguistic meaning of the source document in the target language. But the roles played by its various components, natural as well as artificial, are far from trivial, and the division of linguistic labor between them throws new light on the familiar notions that were initially inspired by rather different phenomena in the philosophy of language, mind, and cognitive science.

Keywords Extended mind · Computer-assisted translation · Machine translation · Translation memory

✉ Yuri Balashov
yuri@uga.edu

1 Department of Philosophy, University of Georgia, Athens, GA 30605, USA

1 Introduction

Human translation is a rapidly growing industry.[1] Professional translators specializing in technical areas,[2] ranging from aerospace engineering and medical equipment to financial management and international law, pride themselves on the quality of their work and its indispensable role in fostering intercultural communication, facilitating global business development, and contributing to the progress of science and technology around the world. For these and other reasons, translators view their work as a valuable service.

However, for nearly 70 years human translation has been facing competition from machine translation (MT). Since its emergence in a highly polarized post-war environment of the early 1950s, the primary explicit goal of MT has been to make human translators irrelevant by rendering their services unnecessary. This goal has not been—and may never be—realized. But it continues to fuel the rapid development of MT, which has gone through several stages.[3] There is no question that MT has come a long way and its quality has greatly improved. However, contrary to many bold claims,[4] this has not resulted in making human translators redundant. It did result in rethinking the entire relationship between human and machine translation. While this relationship has historically been adversarial, an increasing number of stakeholders on both sides are now coming to view it as symbiotic instead.

The present case study aims to explore this new understanding. This is not possible without taking a close look at how professional human translators do their work, especially since it has changed rather dramatically in the last three decades due to a combination of factors: the development of natural language processing techniques, the increased capabilities of personal computers, and the rapid growth of the amount of mono- and bilingual data available to human translators at their fingertips. Nowadays most translators perform their work in a computer-controlled translation environment using CAT/TM tools.[5]

The evolution of these tools is an underexplored topic in the history of science and technology.[6] But it has been overshadowed by the much more spectacular and

[1] The U.S. Bureau of Labor Statistics (https://www.bls.gov/ooh/media-and-communication/interpreters-and-translators.htm, Apr 28, 2020) predicts a 19% growth in translation job opportunities between 2018 and 2028, which is much higher than the 5% average growth for all careers.

[2] From now on, by *translation* I will mean *technical* (*non-literary*) "written translation," as opposed to "oral translation" known as *interpreting*.

[3] These stages and their approximate timelines are as follows: *rule-based* MT (1950s–1990s), *statistical* MT (1990s–2015), and *neural* MT (2015–present). On the history of rule-based MT, see Hutchins (1986) and Hutchins and Somers (1992). For an authoritative introduction to statistical MT, see Koehn (2010). For a technical review of both approaches as of 2008, see Jurafsky and Martin (2008: Ch. 25). Both Koehn (2010) and Jurafsky and Martin (2008) briefly discuss the history of the field. Poibeau (2017) is a popular account that just barely touches the beginnings of neural MT. Koehn (2020) is a state of the art introduction to neural MT. The history of MT lies beyond the scope of the present paper.

[4] See, in particular, Wu et al. (2016) and Hassan et al. (2018).

[5] 'CAT' stands for *computer-assisted* (or *-aided*) *translation*, 'TM' for *translation memory*. In speaking of tools, I will use these terms interchangeably.

[6] See Hutchins (1998), Somers (2003), and Sin-wai (2017).

sensational development of machine translation systems, with which TM tools are sometimes confused (with the terminology—'TM' vs. 'MT'—adding to the confusion). Despite all the differences between these two language technologies there are some similarities in the algorithms they use. Even more importantly, these technologies now interact in mutually enhancing ways, and their increasing role in human translation has become a subject of ergonomic and process studies.[7] Philosophers and linguists, however, have been slow in coming to grips with these significant developments in the interrelated areas of human and machine translation.

The present paper seeks to fill in this lacuna by exploring in some detail the highly distributed nature of the human–computer interaction in the CAT process, which presents a remarkable case of an *extended cognitive system* involving a translator, a TM tool, an MT engine, and sometimes other human translators or editors. Considered as a whole, such a system is in the business of representing the *linguistic meaning* of the source document in the target language. But the roles played by its various components, natural as well as artificial, are far from trivial, and the division of linguistic labor between them throws new light on the familiar notions that were initially inspired by rather different phenomena in the philosophy of language, philosophy of mind, and cognitive science.[8] Language translation gets at the very heart of meaning representation in human, computer, and hybrid human–computer systems.[9] A close look at how these two components of the process interact in real life may offer new insights into how physical systems represent linguistic meaning.

To place these issues in their proper context I start with an introduction to CAT tools and their central functions (Sects. 2 and 3). As we go along, I flag several phenomena for special consideration and discuss them from a more theoretical angle in Sects. 4 and 5, focusing on the role of semantic substitutivity in *fuzzy match repair* and *smart fragment assembly* which, I argue, demonstrate the distinctive features of the translator's *extended mind*. I also draw a contrast between the latter and a popular response to the Chinese room argument. In Sect. 6 I turn to more recent developments in TM-MT symbiosis: *interactive translation prediction* and *adaptive machine translation*, which provide striking examples of *sub-symbolic predictive processing* at work. Finally, in Sect. 7 I review recent advances in *cognitive translation studies* and discuss potential new avenues of interdisciplinary research into the nature of the translator's extended mind.

[7] See, e.g., Sánchez-Gijón et al. (2019); Daems and Macken (2019); Knowles, Sanchez-Torron and Koehn (2019).

[8] For a systematic discussion of the extended mind debates by a leading proponent of the idea, see Clark (2008).

[9] "Anyone who seriously tries to understand what translation is about must come close enough to the philosophical disputes about meaning to feel the heat, if not to see the light," notes Martin Kay (2017: 49), a prominent computational linguist who greatly contributed to the development of computer tools for human translators early in the process (Kay 1980).

2 CAT Tools and Their Functions

Human translators working on technical projects often encounter similar documents; for example, equipment manuals changing little from year to year and between models, or prescription drug information which has mandatory section headings and standard language inside the sections, which varies from document to document mostly in the names of substances and numerical values of various parameters. *Translation memory* (TM) is a way to leverage such similarities.

The idea of storing human translations in a growing database and querying them when a similar sentence is encountered goes back at least to the early 1980s.[10] The implementation of the idea, however, had to wait until the 1990s when substantial progress was made in generally available word-processing technologies, and personal computers with sufficient CPU and RAM resources appeared on the market.

A contemporary CAT/TM tool[11] divides the source text (the text to be translated) into segments (typically sentences), allowing the translator to work on them one by one, and enters human-confirmed translations into a translation memory. When a new segment is opened for translation the software looks for *exact* and *fuzzy matches* in the TM and presents them to the translator for consideration, modification, and endorsement. In addition to sparing the user from duplicating the work already done and thus boosting productivity, the system creates a friendly environment allowing the translator to focus on a given source sentence and its immediate context and enforces consistency of translation within and across the documents.

Besides TMs containing sentence-long bilingual units, most CAT tools also provide for *sub-sentential terminology management* allowing one to create, modify, look up, and reuse the translations of individual domain-specific terms stored in *term bases* (TB) or specialized glossaries.[12] The interplay between these two types of bilingual resources, sentential and sub-sentential, is key to the proper functioning of the translator's "extended mind," as I argue in Sects. 3 and 5. The present section sets the stage for this analysis by illustrating some basic features of modern CAT tools. In my illustration I use memoQ Translator Pro 9.1.8 and the French → English language pair. The examples below are based on the French–English part of the EMEA multilingual parallel corpus constructed from the documents of the European Medicines Agency.[13]

[10] See Kay (1980), Hutchins (1998), Somers (2003), and Sin-wai (2017).

[11] Such as SDL Trados Studio (https://www.sdltrados.com), memoQ (https://www.memoq.com), Wordfast (htpps://www.wordfast.com), Déjà Vu (https://atril.com), OmegaT (https://omegat.org), and many others.

[12] In the simplest case, a term base is a two-column electronic table matching specialized source language words (e.g. *tubulaire*) and multiword expressions (such as *muscle cardiaque*) with their target language counterparts (*tubular* and *myocard*, respectively).

[13] opus.nlpl.eu/EMEA.php. See Tiedemann (2012).

2.1 Computer-assisted Translation Environment

The actual work in a CAT tool is done in the *translation grid* (Fig. 1), which displays the segments with source on the left and target on the right allowing the translator to navigate through the document and to enter and revise the translations of individual segments. The *translation results* pane on the right side of the editor window displays color-coded translation memory matches.[14] The TM suggestion displayed in Fig. 1 has the fuzzy match rate of 77%.[15] The tracked changes in the lower portion of the translation results show the difference between the source text and the TM source language entry.[16]

The translator can work on the source document segments in any order. Once the final translation for a segment is confirmed it is entered into the working TM and the tool moves to the next segment. TB entries are usefully highlighted in the source text, and their translations can be inserted at the cursor point in the target segment with a single hot key or mouse click. Term translations are also suggested as floating predictions once their first character or two are entered.

2.2 Concordance Search and Predictive Typing

TM fuzzy matches below 75% tend not to be very helpful. For this reason the match threshold is often set at that level. This means that there will be no TM matches for many sentences such as one below:

> *Replagal est une solution pour perfusion dont la substance active est l'agalsidase alfa.*

But parts of such subthreshold sentences may still be found in the TMs. *Concordance search* allows one to retrieve these parts from their surrounding linguistic context and use them in assembling a translation. Below are some concordance search results for the sub-sentential string *est une solution*, along with the translation for the first one retrieved by the CAT tool from the TM:

> …
> *Pedea* ***est une solution*** *transparente injectable conditionnée dans une ampoule.*
> *[Pedea is a clear solution in an ampoule for injection.]*
> *Le solvant* ***est une solution*** *liquide incolore.*
> *Prialt* ***est une solution*** *pour perfusion contenant le principe actif ziconotide…*
> …

---

[14] As well as other assorted suggestions coming from different sources, such as *longest substring concordance*, *fragment assembly*, and *machine translation plug-ins*. More on them below.

[15] Fuzzy match rate is a measure of similarity between the source text and the automatically retrieved TM source language entry.

[16] A disclaimer: some of the EMEA corpus segments used in illustrations here and below were deliberately changed from their original form for demonstration purposes and should not be used for any other purpose.

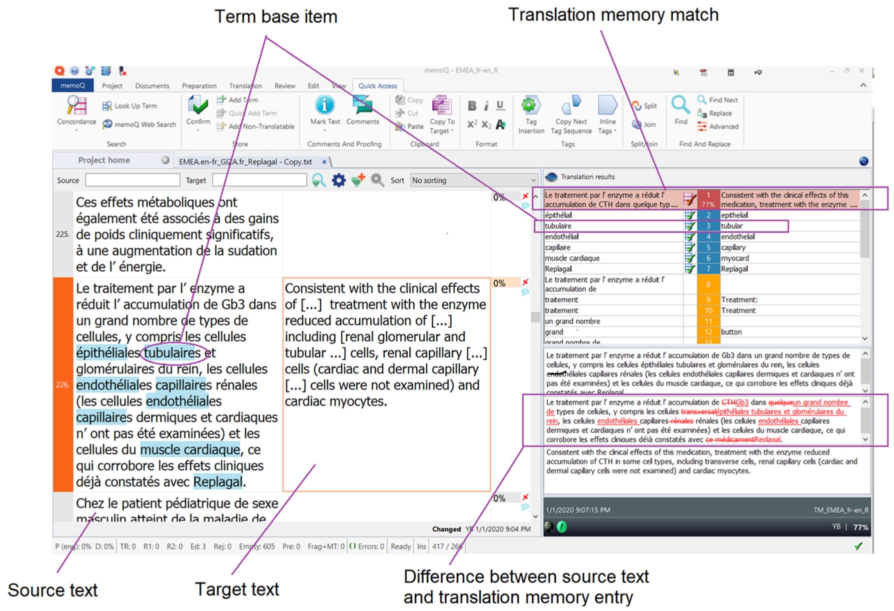

Fig. 1 Translation environment in memoQ. The *translation grid* on the left displays the source and target language segments (sentences) side by side, allowing the translator to navigate through the document and to enter and revise the translations of individual segments. The *translation results* pane on the right side presents color-coded translation memory matches and glossary (term base) items for an active segment. The tracked changes in the lower portion of the translation results show the difference between the source text and the TM source language entry

Performing too many concordance searches is impractical and time consuming; but automatic identification of the *longest* concordance substrings may be very helpful. The "orange" items in the translation results (items 8-12 in Fig. 1) come from *longest substring concordance*. In some cases a CAT tool can offer a translation of an identified substring, as shown in Fig. 2 below. This happens when the substring also occurs as a stand-alone source-language unit in the translation memory.

Another useful feature of CAT tools has to do with *predictive processing*.[17] The algorithm behind it is similar to those at work in *statistical machine translation*. In memoQ this feature is called a *Muse*:

> A Muse is a resource—a statistical database—that offers hints for predictive typing…. A Muse guesses the next word or expression, computing from the source text and the translation that was already typed. In a way, the Muse is

[17] Which was already mentioned in connection with specialized term insertion. Predictive processing is not just a time-saving trick; it is a cognitively important phenomenon which has received much attention lately (see Hohwy 2013; Clark 2014: Ch. 11). It lies at the basis of interactive and adaptive human–machine translation, the topic of section Sect. 6.2.

Fig. 2 Automatic translation suggestion for a longest concordance substring (…*solution pour perfusion*…) as it occurs in *Replagal est une solution pour perfusion dont la substance active est l'agalsidase alfa*

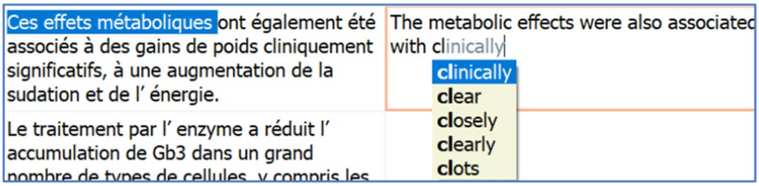

Fig. 3 Predictive typing suggestions from a memoQ Muse

> similar to a statistical machine translation module, but it can be more precise—because it does not have to guess an entire sentence, just the next few words.[18]

A Muse can be trained on any number of TMs and other bilingual resources. Its predictive typing feature is illustrated in Fig. 3.

In addition to the functional features and operation modes mentioned above, modern CAT tools have many others which are too numerous to describe here. But what has been said gives a good idea of the enormous gains in efficiency, consistency, and precision afforded by these tools to human translators and provides enough background to illustrate the highly distributed and symbiotic nature of the human–computer translation process in the following section. This is best done by walking through a real example.

3 Human–Computer Interaction in CAT Workflow

Put yourself in the shoes of a typical medical translator working on a French-into-English pharmaceutical project (Fig. 1). How would she go about translating the following segment? (Fig. 4)

We can safely assume that the translator's command of the source language (French) is very good, that she is translating into her mother tongue (English) and

[18] https://docs-memoq-com.azurewebsites.net/current/en/Places/project-home-muses.html. CAT tool developers do not reveal the details of their algorithms. But it is safe to assume that memoQ's Muses and similar auto-suggest functions of other CAT tools are a far cry from the state-of-the-art contextual language generation systems based on neural networks. Since the latter have demonstrated much better text prediction performance than the classical *n*-gram language models (used in statistical MT, which dominated the field before the advent of neural MT around 2016), CAT tools might benefit from incorporating more advanced AI-driven algorithms into their bilingual search and prediction functions. Unfortunately, this may be difficult to implement in traditional CAT tools on a mass scale (see notes 20 and 25). But recent successful deployment of *interactive translation prediction* and *adaptive machine translation* in the framework of "smart CAT tools" appears to be very promising. I return to these developments in Sect. 6.2.

Le traitement par l'enzyme a réduit l'accumulation de Gb3 dans un grand nombre de types de cellules, y compris les cellules épithéliales tubulaires et glomérulaires du rein, les cellules endothéliales capillaires rénales (les cellules endothéliales capillaires dermiques et cardiaques n'ont pas été examinées) et les cellules du muscle cardiaque, ce qui corrobore les effets cliniques déjà constatés avec Replagal.

Fig. 4 Source segment No. 226 from Fig. 1

FR: Le traitement par l'enzyme a réduit l'accumulation de ~~CTH~~Gb3 dans ~~quelques~~un grand nombre de types de cellules, y compris les cellules ~~transversal~~épithéliales tubulaires et glomérulaires du rein, les cellules endothéliales capillaires rénales (les cellules endothéliales capillaires dermiques et cardiaques n'ont pas été examinées) et les cellules du muscle cardiaque, ce qui corrobore les effets cliniques déjà constatés avec ~~ce médicament~~Replagal.

EN: Consistent with the clinical effects of this medication, treatment with the enzyme reduced accumulation of CTH in some cell types, including transverse cells, renal capillary cells (cardiac and dermal capillary cells were not examined) and cardiac myocytes.

Fig. 5 Translation memory match (77%) for the source sentence shown in Fig. 4. The difference between the TM match and the source sentence is shown as tracked changes

is generally familiar with medical terminology in her language pair. She knows, in particular, that *rein* means *kidney* and *tubulaires* means *tubular*, and that *cellule* is normally translated in medical contexts as *cell* (rather than *cage* or *tank*) and *cellule du muscle cardiaque* as *cardiac myocyte* (rather than *heart muscle cell*). However, she has no detailed knowledge of renal cytoarchitecture. Specifically, she has never seen the images of renal cells, does not know how *les cellules endothéliales capillaires dermiques* and *les cellules endothéliales cardiaques* differ from each other, and so forth. This is a very typical situation.

Fortunately, she has a good translation memory and a good term base specialized to the area in question, and she can fully trust these bilingual resources. In fact, her TM offers a 77% match for the above sentence (see Fig. 5).

Additionally, the term base for the project has entries automatically highlighted for the translator's attention in the source segment (see Fig. 1).

These resources can be leveraged in somewhat different ways depending on the extent of the translator's knowledge of the domain in question (in our case, it's rather moderate), the proximity of the source and target languages (rather close), the register and style of the document to be translated, the specific requirements and guidelines, and other parameters of the project. One possible strategy is described below, step by step.

3.1 Deciding on the Overall Syntactic Frame of the Target Sentence

Considering the grammatical similarities between English and French it may be tempting to simply follow the order of the source sentence which has a relative clause at the end. In that case the translation would begin with *Treatment with the* *enzyme…* and end with *…which corroborates the clinical effects previously seen with Replagal*. The TM, however, suggests a different syntactic form for the translation where the relative clause is *fronted*. While this way of putting it (*Consistent with the clinical effects of X…*) may sound unnatural to the non-specialized ear it may be standard language in the pharmaceutical documentation. Usefully reflected in the TM, this fact about usage puts the translator on a good course.

3.2 Span Pre-translation

Once this syntactic decision is made the translator can look carefully at the differences between the TM match and the source sentence. Color coding and red line tracking (Figs. 1 and 5), which show the differences, play an important role in focusing the translator's visual attention on the relevant fragments of the source sentence. The next natural step is to determine *which target words* in the above TM unit are affected by the mismatch. This is done by mentally *aligning* the tracked source-language part of this unit with its English counterpart and mentally *locking* the *pre-translated spans* of the text (shaded in Fig. 6), which will prevent them from further change and allow the translator, in effect, to work around them.[19]

In general, word and phrase alignment of this kind requires a combination of morphological, syntactic, and semantic knowledge with "permutation" skills in both languages in order to enable the translator to move the relevant parts of the target sentence across its frame in her head while keeping track of their dependency relations and other grammatical roles. However, this knowledge need not be perfect, and the required permutations need not rely on a full syntactic analysis of the source sentence. As already noted, the translator may lack expert knowledge of the special terms figuring in the source sentence. Instead she can rely on their term base translations. And she can mentally manipulate entire phrases such as *Consistent with the clinical effects of...* as integrated syntactico-semantic units without resolving them into their ultimate constituents. Multiple movements of this sort, however, put a burden on working memory, especially when dealing with morpho-syntactically different languages and longer sentences. The task is further exacerbated by the fact that many aligned spans are neither grammatical phrases nor even units that could draw the translator's attention on their own. The translator is forced to deal with "linguistically unaware" products of the fuzzy-matching algorithm—such as *cell types, including*—if she wants to make full use of her TM.[20] So in addition to

[19] *Span pre-translation* is a term that was first introduced to describe a somewhat similar process in machine not human translation (see Vandeghinste et al. 2017).

[20] Could fuzzy matches be made more "linguistically aware"? This is a translator's dream, one among many, that unfortunately cannot be realized without enhancing CAT tools with rather advanced AI features, which is impossible to implement in a uniform and economically reasonable way for over 100 morphologically and syntactically diverse languages supported by the tools. Despite their indispensable role in extending the translator's mind, CAT tools are, almost paradoxically, *language-blind*. Exploring more advanced fuzzy match options for a given language pair in a restricted framework of a research project is, of course, a different matter. For interesting recent work along these lines on *linguistically aware* fuzzy matches, see Vanallemeersch and Vandeghinste (2015).

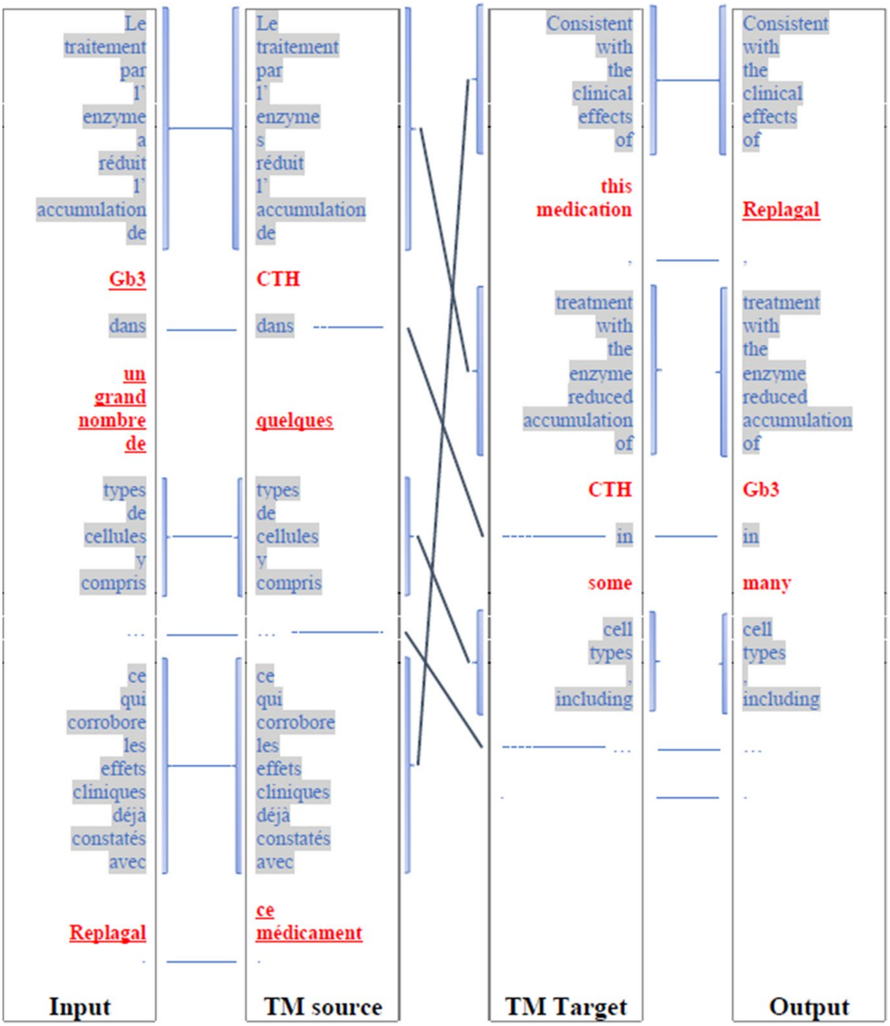

Fig. 6 Span pre-translation; modeled after Vandeghinste et al. 2017. Which target words are affected by the mismatch between the input sentence and the translation memory (TM) source entry? The shaded portions of the input sentence represent pre-translated spans already contained in the TM. They are aligned (boldfaced lines in the middle) with their target counterparts. Note that the TM suggests moving the relative clause at the end of the source sentence to the front of the target sentence. The red/boldfaced parts of the input sentence and the TM source represent the differences between them. The translator can "lock" the pre-translated spans of the input text in her head and "word around them" translating only the unaligned red/boldface words and phrases

mentally aligning the corresponding spans of the source and target text the translator must also be able to "interpolate" across their boundaries in order to anchor source and target span correspondences in the underlying syntactic and semantic relations which have their own structure and logic. At some point there may be too many things to keep in the air. The translator can alleviate this cognitive burden quite significantly if she performs the required manipulations sequentially in the translation grid. Indeed, visualizing some of these steps by typing them into the target segment not only reduces the memory load but sometimes leads to further linguistic insights that may be unavailable when the manipulations are performed in the head.[21]

Span pre-translation, in turn, delineates a sequence of local tasks involving word and short-phrase substitutions, deletions, and insertions in the target language intended to reproduce the corresponding patterns in the source segment. Let us go over these operations for our sample sentence.

3.3 Local Semantic Interpolation

The first operation involves the transition from *…l'accumulation de **CTH** dans…* to *…l'accumulation de **Gb3** dans…* in the source sentence, with *Gb3* coming to replace *CTH*. What do these terms mean? The translator may not know or remember; and her project resources are of no help in this case: TB and TM concordance searches return no occurrences of either term. The translator must resort to an external search, which quickly reveals[22] that *Gb3* and *CTH* are *synonymous* international acronyms for a lipid variously known as *globotriaosylcéramide* (*globotriaosylceramide)* or *céramide trihexoside* (*ceramide trihexoside*). Thus the replacement does not affect the denotation of the term; but its alternate versions may still be differentially appropriate for different contexts. Synonymy is an important semantic relation. Accordingly, the equivalence of the two terms (and of their fully spelled out versions) may be worth recording for the future. This can be done by adding new entries to the term base, which takes a bit of time and attention, and some extra eye and finger movement. But the effort may pay off later.

Whereas the second local change — *…dans **quelques** types de cellules…* → *…dans **un grand nombre de** types de cellules…* —is straightforward, the next one involves several term replacements and insertions that require more work. There *les cellules **transversal*** gives way to *les cellules **épithéliales tubulaires et glomérulaires du rein***, and ***endothéliales*** gets inserted on two occasions. While most of the special terms occurring in this passage are conveniently attested in the term base (and highlighted in the source column in Fig. 1) the substitutions in question require syntactic and semantic decisions that must be informed not only by knowledge of the French and English grammar[23] but also by knowledge of the world. Should *les cellules épithéliales tubulaires et glomérulaires du rein* be parsed as *les cellules [[épithéliales tubulaires] et [glomérulaires du rein]]* (and translated as *tubular epithelial and renal glomerular cells*) or rather as *les cellules [[[épithéliales tubulaires]*

---

[21] Such positive feedback between "external" visual representation and "internal" mental linguistic representation is typical of cognitive extensions of the human mind (see, e.g., Clark 2008 and 2014: Ch. 9). We will revisit it in the context of interactive/adaptive human–machine translation in Sect. 6.2.

[22] See e.g. https://www.ncbi.nlm.nih.gov/pubmed/17073606 (Feb 2, 2020).

[23] While both languages are mostly head-initial they diverge when it comes to noun phrases; cf. *les cellules transversal* vs. *transverse cells*.

et glomérulaires] du rein] (*renal tubular epithelial and glomerular cells*)? Here some knowledge of cytology might help (both types of cells are renal cells), and the immediate context suggests this reading too. The risk of mis-interpretation can be further reduced by googling *renal tubular epithelial cells* and *renal glomerular cells* as exact quoted expressions.

The same dilemma arises for *les cellules [[endothéliales capillaires dermiques] et [cardiaques]]* vs. *les cellules [[endothéliales capillaires] [dermiques et cardiaques]]*. Should this phrase be read (in English) as *dermal capillary endothelial and cardiac cells* or *dermal and cardiac capillary endothelial cells*? Here all the resources available to the translator are insufficient to resolve the ambiguity of the French phrase. Unless the translator has an expert insight into the nature of these cells in the context of treatment with Replagal she appears to be stuck. But there is a work-around in this case: *reproduce* the ambiguity of the French phrase in English by switching the order of *dermal* and *cardiac*. The result—*cardiac and dermal capillary endothelial cells*—is consistent with both readings and rides on a parameter setting difference between French and English noun phrases. Faithful reproduction of linguistic ambiguities is fair game in translation. Notice that this strategy is unavailable in the first case considered a paragraph ago.

After the last change, which is straightforward, — …*avec* ***ce médicament***→ … *avec* ***Replagal***—the entire English translation can finally be assembled, proofread, endorsed, and committed to the translation memory (Fig. 7 below).

The above account ignores other challenges that often arise in the translation process, for example *co-reference resolution*. But hopefully, it does demonstrate that human translation, even when augmented with computer tools, is indeed a very intensive cognitive process which requires a broad range of linguistic and non-linguistics skills. These range from search techniques, memory load control, multitasking, rapid mental set switching, and other executive functions to the abilities to survey and keep in the air many syntactic, morphological, semantic, and pragmatic parameters of both languages, which may be interrelated in complicated ways. The translator must reconcile their often-conflicting grammatical demands and resort to creative ways of handling linguistic ambiguity, among other things.

> Consistent with the clinical effects of Replagal, treatment with the enzyme reduced accumulation of Gb3 in many cell types, including renal glomerular and tubular epithelial cells, renal capillary endothelial cells (cardiac and dermal capillary endothelial cells were not examined) and cardiac myocytes.

Fig. 7 Target segment No. 226 (human translation using a CAT tool)

4 Semantic Substitution, Fuzzy Match Repair, and Smart Fragment Assembly

From a more theoretical viewpoint, the process described above makes good use of *semantic substitution*. Specifically, the "interpolation" between pre-translation spans turns on replacing a one-, multi-, or zero-word[24] phrase α in a source sentence *S* with β:

$$S \rightarrow S[\beta/\alpha]$$

and reproducing the result of such a replacement in the target sentence. Assuming that the two languages are close enough to allow for span pre-translation and interpolation and that Tr(α) and Tr(β) are official term base translations of α and β respectively, the meaning of the translated sentence after the replacement, $\left\|\mathrm{Tr}\ (S[\beta/\alpha])\right\|$, should be related to the meaning of the original sentence $\|\mathrm{Tr}\ (S)\|$ as follows:

$$\left\|\mathrm{Tr}\ (S[\beta/\alpha])\right\| = \left\|\mathrm{Tr}(S)[\mathrm{Tr}(\beta)/\mathrm{Tr}(\alpha)]\right\|$$

This familiar systematic connection between the meaning of sentential and sub-sentential expressions in both languages is precisely what allows the translator to perform the combinatorial tasks described above and to offer a good translation of the source sentence even if she does not know the exact meaning of certain sub-sentential expressions. Instead, she can rely on their translations retrieved from the translation memory and the term base.

Semantic manipulations of this sort can be automated. Many contemporary CAT tools offer a feature known as *fuzzy match repair* (FMR). It is based on structured cooperation between translation memories and term bases which often happens behind the scenes and makes use of essentially the same mechanics of semantic substitution. If a source sentence differs from a TM match in a single term *and* if both terms also figure in a TB, as in the example below, this gives the system all the resources needed to generate the correct translation of the source sentence, even though that translation is not present in the TM (Fig. 8).

Here *S* = *Replagal est recommandé pour un usage à long terme*, Tr (*S*) = *Replagal is intended for short-term use*, α = *court terme*, β = *long terme*, Tr (α) = *short-term*, and Tr (β) = *long-term*. Consequently, Tr (*S* [β/α]) = Tr (*S*) [Tr (β)/Tr (α)] = *Replagal is intended for long-term use*, which a CAT tool can offer as a "patched" translation of *S* [β/α]. In many cases this can be accomplished with little or no input from the translator. All it takes is a relatively simple, language-blind string search and replacement algorithm complete with regular expressions. Numerical and other mismatches between "untranslatables" (dates, times, medication and geographic names, etc.) can be repaired in a similar way but without any aid from term bases.

[24] Null strings are usually notated as ε. α = ε corresponds to an insertion and β = ε to a deletion.

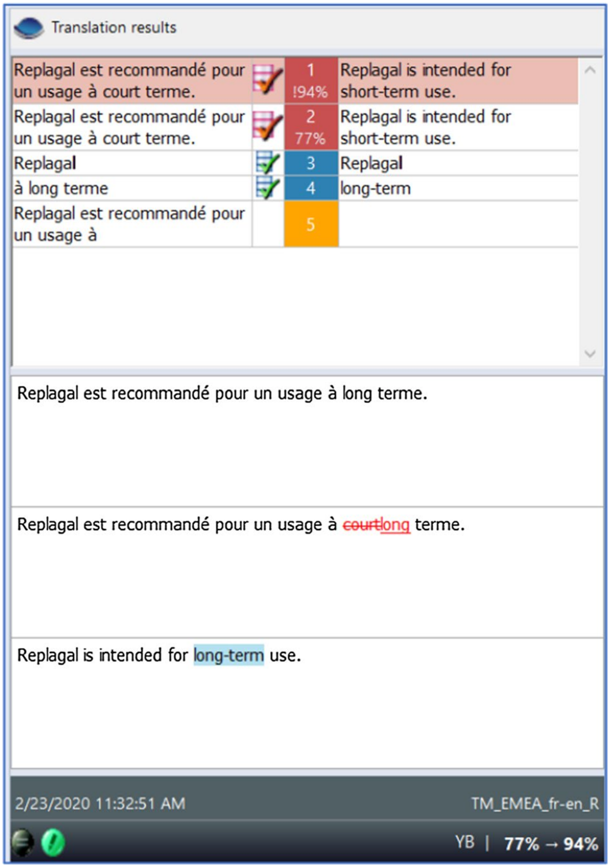

Fig. 8 A fuzzy match repair ("match patch") in memoQ

Now let us turn to a more sophisticated case of computational semantics at work in computer-assisted human translation. Consider this sentence:

Source segment No. 252:
Les données non cliniques issues des études conventionnelles de toxicologie en administration répétée n'ont pas révélé de risque particulier pour l'homme.

for which the TM provides several fuzzy matches ranging from 67% to 79% (Fig. 9a). What is especially notable in the present context is that the tool identified three non-overlapping longest substrings (marked '6', '7', and '8' in Fig. 9a) which, when concatenated, comprise the entire source sentence $S_{252} = \alpha_6\alpha_7\alpha_8 =$ *Les données non cliniques issues des études conventionnelles de toxicologie en administration répétée n'ont pas révélé de risque particulier pour l'homme*. The tool does not offer their translations, but they can be easily retrieved from three concordance search results. For example, the first set of results allows one to glean the translation of

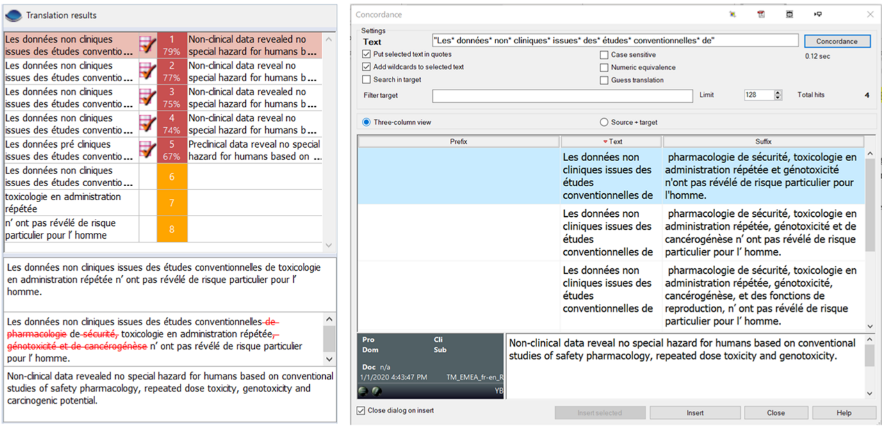

(a) (b)

Fig. 9 **a** Three non-overlapping longest concordance substrings ('6','7', and '8'), identified by the CAT tool, comprise the entire source sentence. **b** Some concordance search results for substring '6' retrieved from the translation memory

α_6 = *Les données non cliniques issues des études conventionnelles de* with minimal semantic knowledge by identifying the corresponding common target span in the TM matches displayed at the bottom of the concordance search window (Fig. 9b): Tr (α_6) = *Non-clinical data […] based on conventional studies of*. The translations of the other two substrings can be mined from the concordance data in the same way. Putting all three together (and moving *no special hazard for humans* to *[…]*, as suggested by the TM) produces the translation of the entire sentence, Tr (S_{252}) = *Non-clinical data reveal no special hazard for humans based on conventional studies of repeated dose toxicity*, whose meaning is determined as follows:

$$\left\|\mathrm{Tr}\ (\alpha_6\alpha_7\alpha_8)\right\| = \left\|\mathrm{Tr}\ (\alpha_6)\mathrm{Tr}(\alpha_7)\mathrm{Tr}(\alpha_8)\right\|$$

Two brief comments are in order. (1) There is no reason why "smart" fragment assembly of this kind cannot be performed automatically other than the practical computational limitations deliberately imposed by the developers of contemporary CAT tools on their products.[25] (2) The source sentence considered above is quite

[25] (a) CAT tools must run very fast on generally affordable desktop or laptop computers. (b) As already mentioned, CAT tools are *language-blind*: they apply the same "linguistically unaware" algorithms to over 10,000 supported language pairs, despite the numerous grammatical differences between them. But: (i) the computational power of personal computers has been steadily increasing and its cost rapidly decreasing in the recent years; (ii) larger language service providers are actively working to overcome the algorithmic limitations of CAT tools by customizing them to specific language pairs and integrating them with machine translation; and (iii) some CAT tools already accomplish automatic smart fragment assembly in some cases.

exceptional in that it can be broken into three non-overlapping substrings which are, furthermore, (almost) complete grammatical phrases that combine with each other in familiar syntactic ways providing for a rather straightforward semantic analysis. This does not happen very often, even for languages as close as French and English.

Still, the phenomena in question (i.e. automatic fuzzy match repair and smart fragment assembly) are notable as they demonstrate already implemented or technologically possible ways in which linguistic labor can be shared between a human translator and a CAT tool. This division of linguistic labor calls for further analysis.

5 The Chinese Room of Computer-Assisted Human Translation?

The description of a typical CAT process in the previous section makes it clear that both key participants in it—the human translator and the TM tool—contribute to the process in symbiotic ways, hammering out the target sentence in stages involving trial and error, search, replacements, deletions and insertions, mutual adjustment (morphological, syntactic, semantic), and other actions, some of which can be characterized as decidedly *epistemic*.[26] The translator is often in the driver's seat.[27] But in principle, he need not always be there, because he can outsource important semantic responsibilities to his entrusted bilingual resources—translation memories (TM) and term bases (TB).[28]

High-quality TMs and TBs are painstakingly developed and maintained by translators over many years and are their most valuable assets. The majority of the TM and TB segments are generated during actual work in a given specialty field performed by one or more translators, others may come from alignment of existing parallel documents (e.g. published pharmaceutical leaflets, as in our illustration above) produced by human translators in the past. Typical TMs have tens of thousands translation units (i.e. pairs of source and target sentences), and TBs may have hundreds of terms entries. This makes them humanly unsurveyable. No translator is under obligation to internalize their content. Indeed, CAT tools were designed precisely to relieve him of such excessive cognitive burden. Furthermore, the translator need not be an expert in the technical terminology residing in these large repositories of bilingual data (Sect. 3). When a combined TM-TB system suggests term Tr (β) as a replacement of Tr (α) in a fuzzy match repair the translator may not

[26] In the sense relevant to cognitive science, *epistemic actions* are "physical actions that make mental computation easier, faster or more reliable"; they are "designed to change the input to an agent's information-processing system" by "modifying the external environment [in order] to provide crucial bits of information just when they are needed most" (Kirsh and Maglio 1994; quoted in Clark 2014: 194).

[27] E.g. when it comes to resolving co-reference or word sense disambiguation based on broader knowledge of the world.

[28] The emphasis on the *resources* is important. It has been repeatedly noted above that CAT tools are language-blind. If so, how can they discharge any *semantic* responsibilities? The answer turns on drawing a distinction between CAT *software*, which are indeed language-blind, and bilingual *resources* (i.e. TMs and TBs), which are not. The software can assume the requisite responsibilities by *using* the resources as *data*.

be aware of their precise technical meaning. The same applies to longer substrings when a CAT tool identifies them, and sometimes puts together, offering the result to the human translator to finalize and confirm. The translator remains largely responsible for ensuring the grammatical integrity of the resulting sentences—something CAT tools may be unable to do on their own in the absence of more advanced AI-based methods. But the required morpho-syntactic polishing can be performed in the absence of deep semantic knowledge; all one needs is knowledge of part-of-speech identity of the relevant unfamiliar terms and their shallow semantic features such as animate vs. inanimate, solid vs. liquid, and the like. As a result, the translator's semantic obligations may be rather limited. To perform quality work, he does need to have some familiarity with the broad area in which he is working; indeed, specialization is a must in the translation profession. But as already noted, he cannot be expected to be an expert in a narrow technical subfield. For example, the translator working on medical equipment should be familiar with the bilingual technical terminology in this general area; but he can hardly be required to know the distinction between different types of intra-aortic catheters or have a clear mental image of them. The translator must have excellent search skills allowing him to establish term and phrase equivalences in such cases beyond reasonable doubt without becoming an expert on intra-aortic catheters. In doing so, he relies on the above-described features of his CAT tool (i.e. fuzzy match repair and smart fragment assembly).[29]

So in the end, the translator can sign off on a morpho-syntactically coherent target language sentence whose exact technical meaning is beyond his grasp. Instead, he can delegate considerable semantic responsibility to his resources—the translation memory and the term base—which serve as huge containers of specialized knowledge distributed over many overlapping translation units and term pairs. This, of course, assumes that the resources are good and can be trusted. But the situation seems to be no different from Putnam's experts on gold whose professionalism is needed to underwrite ordinary folk's talk of gold and their extra-linguistic practices such as selling, buying, and wearing gold (Putnam 1975). More generally, the situation appears to be similar to the familiar cases of *active externalism* in which we consult various external resources—textbooks, notebooks, encyclopedia, and their electronic versions—in our everyday work to the extent that they become (almost) an integral part of our mind.[30]

But there are at least two important respects in which a "strongly coupled" system consisting of a human translator and his CAT tool is different from other canonical coupled systems discussed in the extended mind literature. First, the external tools considered in the latter serve as subsidiary providers of information that is then *internalized* and used by the agent to perform a cognitively demanding task: for example, complete a multistage financial calculation or plan a complicated itinerary. In contrast, the translator is not obligated to fully internalize the output of a CAT tool during a fuzzy match repair or fragment assembly. And his own subsequent morpho-syntactic polishing need not cut very deep into the semantic flesh of the target sentence. As

[29] In addition, he has various external sources at his fingertips, such as bilingual electronic dictionaries, online encyclopedias, and translation forums.

[30] See Clark and Chalmers (1998) and Clark (2008).

emphasized above the translator can remain ignorant of the exact meaning of the technical expressions used in such operations.

The second distinction has to do with the fundamental nature of translation work. In his discussion of semantic externalism, Gareth Evans notes that "we constantly use general terms [such as 'microbiologist', 'chlorine' (the stuff in swimming pools), and 'nicotine' (the stuff in cigarettes)] of whose satisfaction conditions we have but the dimmest idea" (1973: 190). This use presupposes the existence of experts in our linguistic community who know what 'nicotine' *really* means. Computer-assisted translation equipped with automated fuzzy match repair, smart fragment assembly and other algorithms, on the other hand, operates at a semantic stage that is *prior* to everyday use (by the target language audience). The translation of a source-language sentence *S* into a target-language sentence Tr (*S*) is supposed to be an answer to the question of what *S really* means. Our discussion of the highly distributed and symbiotic CAT process demonstrates that, unlike Evans's common speaker, the translator can substantially contribute to answering this question without having real expertise in matters expressed by *S*, because he can outsource important semantic tasks to his CAT tool and the available bilingual resources which he trusts and which may originate in the collective work of other translators.

This last point about "collective work" warrants further consideration because it highlights not just the *active* but also the *social* nature of the semantic externalism involved in computer-assisted translation. Many contemporary CAT tools allow a project manager to set up a server- or cloud-based process in which a team of translators can work on the same project remotely (Fig. 10). Among other things, this allows the manager to divide a large document among them and thus expedite the process. Sometimes this is the only way to meet a tight deadline. In this case, the translators use common online resources available to them, but each of them also continuously modifies and extends such resources by entering new translations into a cloud-based TM and occasionally adding new entries to a common TB. One can easily imagine a situation in which Translator A makes good use of a new TM unit created just a second ago by Translator B. This sort of linguistic cooperation, of course, requires the relevant expertise of all the participants. But it also leaves room for consultation, negotiation, and quick exchange of ideas concerning the translation of new sentences and terms. Interestingly, it involves cooperation across time as well

Fig. 10 Cloud-based translation

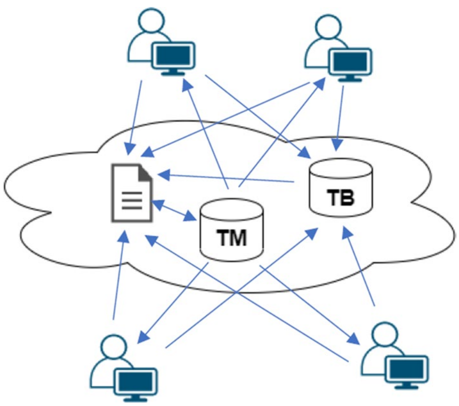

as space. Translators, often residing on different continents, are engaged in online symbiotic workflow with each other; but they also leverage the bilingual data generated by other translators in the past. Their collective activity aimed at representing the meaning of source sentences in a target language may be "smeared quite widely" (Clark 2014: 197) in spacetime across all of them and their common CAT tool.

We have seen earlier that in an individual CAT process the semantic obligations of the translator may be partially delegated to her trusted bilingual resources. This applies to cloud-based projects too, with an important addendum: by interacting with each other in the way described, often in real time, the members of the team can redistribute their semantic responsibilities according to their individual expertise. While no translator can be required to have full knowledge of the domain, each of them can make a valuable contribution to the final product depending on his or her area and level of competence. A target sentence emerging from this synergistic process may include longest substrings grounded in the translation of earlier sentences by Translator A, new terms supplied by Translator B, and so forth. This fresh online delivery of new bilingual data, in turn, enables the CAT tool to perform new fuzzy match repairs and perhaps other more sophisticated operations behind the scenes. Clearly, it is the whole human–computer system, distributed in space and time, that can be said to be responsible for the final translation. The translator's mind is thus extended not only technologically, but also socially.

It may be instructive to compare the picture emerging from this description of contemporary translation work with a popular response to the famous Chinese room thought experiment, known as "the systems reply." In Searle's original scenario (Searle 1980), a monolingual English speaker who knows no Chinese is locked in a room and given a large pile of Chinese writing, and then handed another batch along with extensive instructions in English explaining how to match the (meaningless to him) squiggles in the batch and the pile and manipulate them according to certain English rules, leading him to produce further squiggles and turn them in. Unbeknownst to the English speaker, the pile is a detailed script in Chinese, the batch is a question about it, and his output squiggles are a correct answer to the question. The point of the scenario was to raise an objection to the "strong AI thesis" to the effect that instantiating the right sort of program for purely syntactic symbol manipulation is sufficient for thinking and understanding. The popular reply concedes that the speaker possesses no understanding of the Chinese content; but perhaps the entire system in which he is embedded—including the room, the paper, the pencils, and so on—does. According to Searle, this reply has an air of implausibility: "The idea is that while a person doesn't understand Chinese, somehow the *conjunction* of that person and bits of paper might understand Chinese" (Searle 1980: 421).

Be that as it may, the situation with the translator's "extended mind" appears to be different. As emphasized above, the exact meaning of a target sentence generated in a CAT process may be beyond the grasp of an individual translator. In such a situation, and in light of the considerations adduced above, it seems uncontroversial that only the system as a whole including the CAT tool and perhaps other translators can take joint responsibility for adequately rendering the meaning of the source sentence in the target language. To be sure, the analogy is a bit distant, in two related respects. First, the translator normally remains in full command of the broader context as well as the register and the purpose of the target text. Second, the predicament of the English speaker in the Chinese Room scenario is far more difficult as he does not know a single Chinese word. The translator, in contrast, is expected to have a good command of both languages. But his bilingual semantic knowledge may have substantial gaps that can be usefully filled by other components of the system.

One might point out that there is hardly anything new in this situation. Even the traditional translators used dictionaries—paper dictionaries requiring a laborious lookup process—and other non-electronic reference materials. And they too did not have to be experts in cardiology equipment. What exactly has changed between now and then? One difference is obvious: a contemporary translator is willing to transfer a lot more authority to the external extensions of his mind. Traditional human translators were "micro-managers" in a way their contemporary tech-savvy descendants are not. They are happy to let the computer do as much as is practically possible. But of course, the same happens in many other domains.

What seems unique to translation is the *nature* of the delegated tasks, which have to do with something as deep and sophisticated as *meaning representation*—and not, for example, with budget calculation or itinerary planning. But in the end, isn't meaning representation also a mathematical calculation or computation, just a more complicated one? Perhaps it is "just" a computation. Our brains are, after all, computers. Still, we want to know what kind of computers they are. And how they compute linguistic meaning in a bilingual setting. I will return to this topic in Sect. 7. But first, I want to consider another, more recent and fast-developing form of human–computer interaction in translation.

6 MT and TM: The Rapidly Changing Landscape of Human–Computer Integration

All of the above had to do with (computer-assisted) *human* translation. What about machine translation (MT)? Its impressive progress over the last 70 years has already been mentioned. While it has not resulted (contrary to multiple predictions) in driving human translators out of business, it has a great impact on their everyday work, including its cognitive dimensions.

To set them in perspective, let us start by noting that the quality of raw MT output continues to be very uneven, ranging from excellent to unacceptable. Recall "Source segment No. 226" (Fig. 4) and our long discussion of its translation in a CAT tool guided by a 77% TM match and the need to resolve multiple syntactic ambiguities to produce an acceptable target sentence (Fig. 7). Compare it with the output of Google Translate:

Target segment No. 226 (GoogleMT, Feb 2, 2020):
Enzyme therapy reduced the accumulation of Gb3 in a large number of cell types, including kidney tubular and glomerular epithelial cells, renal capillary endothelial cells (demal and cardiac capillary endothelial cells did not examined) and heart muscle cells, which corroborates the clinical effects already seen with Replagal.

which is neither excellent nor terrible, but clearly unacceptable as the final translation of the target sentence. The MT engine, of course, is unaware of the TM recommendation to front the dependent clause; so it followed the order of the French sentence rather closely. Other issues include a grammatical error, poor choice of target terms ('kidney' vs. 'renal'; 'heart muscle cells' vs. 'cardiac myocytes'), and insensitivity to subtle syntactico-semantic issues arising in translating *les cellules endothéliales capillaires dermiques et cardiaques*.[31]

All things considered, it would be risky for anyone to submit MT-generated translation of an important document to the client. In many cases raw MT output is good enough for *assimilation* or *gisting*—getting the "general drift" of a document in a foreign language, and for certain everyday purposes such as orientation in the local restaurant menus. But it is not advisable to rely on an MT-translated contract or medical guidelines.

This is not to deny the great value of machine translation, far from it, but simply to note that it is still very far from the ideal of "fully automatic high-quality machine translation" proclaimed in the early days of machine translation by its pioneers.[32] What is currently on the agenda is *integration* of human and machine translation. How can it be accomplished? And what implications does this integration have for the larger cognitive, linguistic, and philosophical issues in natural language processing?

6.1 Post-editing of Machine Translation

Historically, the most popular form of such integration has involved *post-editing of machine translation* (PEMT) by human translators. It remains the dominant strategy for many of them, especially since many CAT tools now have plug-ins for machine translation engines[33] which enable the translator to combine automatic population of the target segments with good TM fuzzy matches with getting a raw MT input for the rest of the source document.

PEMT has by now a long history with mixed results and limited enthusiasm among translators. Some MT-generated sentences lend themselves rather easily to post-editing, while others are on the opposite side of the spectrum: trying to post-edit them takes more time and effort than translating from scratch. Sometimes MT suggestions put the translator on a wrong syntactic course; witness the "target segment No. 226" amply considered above. For this and other reasons some translators have a negative attitude toward PEMT, even if they are otherwise open to using MT in their work.

Although balancing the costs and benefits of PEMT often turns on multiple factors and their study has become an industry of its own, a major weakness of this method is that it limits the translation process to just two basic operations: revising good TM fuzzy matches (> 75%) or post-editing MT output for full sentences in all other cases. As already pointed out, TM units are the translator's golden resources,

[31] Cf. our discussion of a creative way to reproduce its ambiguity in Sect. 3.3.

[32] See e.g. Bar-Hillel (1960).

[33] Such as Google Translate, Microsoft Translator, DeepL, KantanMT, and others.

and even relatively short sub-sentential fragments of them are often better in quality than their MT counterparts (see Sects. 2.2 and 4 on longest substring concordance and fragment assembly). It is not a good idea to sacrifice them in favor of substandard MT suggestions simply because the former fall under the match threshold.

In the last decade considerable effort, both in the academic and commercial settings, has gone into exploring other, more sophisticated ways of integrating TM and MT language technologies, sometimes resulting in very impressive prototypes of "smart CAT tools" capable of combining and ranking translation suggestions from multiple sources—MT, TM, and TB—based on new cutting-edge technologies.[34] Two such technologies, *interactive translation prediction* and *adaptive machine translation*, are particularly important from a cognitive science perspective and are briefly considered below.

6.2 Interactive Translation Prediction and Adaptive Machine Translation

Interactive translation prediction (ITP) methods were initially developed in the context of phrase-based statistical machine translation[35] and were motivated by the desire to make the interaction between a human translator and an MT system more structured. Instead of presenting the translator with a full-sentence MT output and letting him post-edit it the system builds the target sentence incrementally, by asking the translator to accept or amend the growing prefix and computing, at each step, the best (or several best) suffix suggestion(s) to extend the prefix. This, in effect, extrapolates the predictive typing model familiar from other applications including the "auto-suggest" or "auto-completion" features of many contemporary CAT tools (see Sect. 2.2 above).

The encoder-decoder model with attention[36] commonly used in neural machine translation (which became the dominant MT technology around 2016) provides an especially natural environment for ITP.[37] This model is implemented in Lilt—an interactive and adaptive MT-TM tool initially designed in the course of academic work at Stanford University.[38]

Lilt is a cloud-based software which allows a user to upload one or more TMs and TBs and import documents for translation. Once uploaded, the TMs and TBs become available for fuzzy matching and terminology lookup. More importantly, they become integrated into Lilt's own centralized *Memory* associated with a given project *and* immediately used for training Lilt's neural machine translation engine (Fig. 11).

Such training is domain-, project-, and user-specific and is in addition to Lilt's own default baseline machine translation resource pre-trained for each language pair. This allows the system to generate MT suggestions for segments that fall below the 75% TM fuzzy match threshold. The translator is of course free to accept or reject

[34] E.g. the SCATE project conducted in 2014–2018 by three Belgian universities (see Vandeghinste et al. 2017, 2019).

[35] First by the developers of TransType (Barrachina et al. 2009), followed by those of CASMACAT (Sanchis-Trilles et al. 2014) and SCATE (Vandeghinste et al. 2019).

[36] Bahdanau et al. (2014).

[37] For details, see Wuebker et al. (2016) and Knowles et al. (2019).

[38] https://www.lilt.com. For details, see Green et al. (2014, 2015).

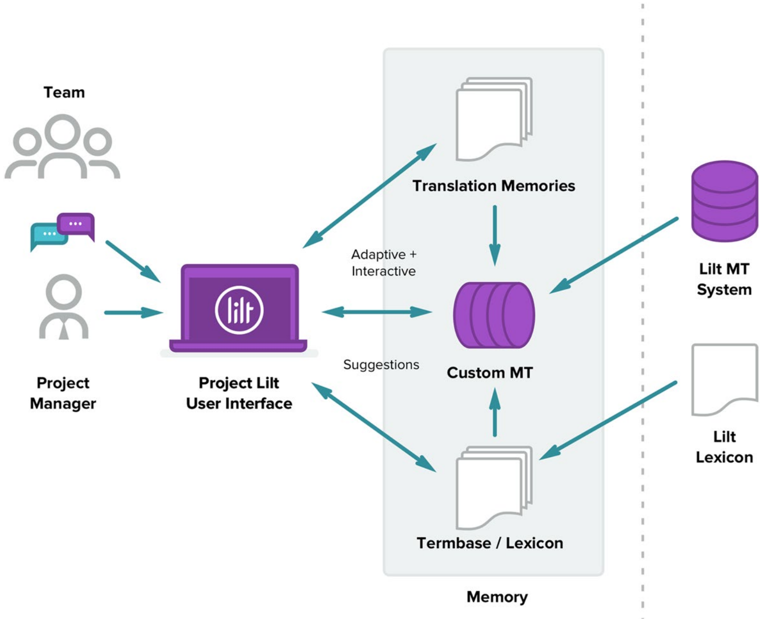

Fig. 11 Lilt's "Memory" (not to be confused with TM) is a collection of source-target sentence pairs which is used "to train the MT system, populate the Translation memory (TM), and update the term base." https://support.lilt.com/hc/en-us/articles/360020608834-What-are-Memories. Accessed Aug 15, 2020

The initial translation suggestion.

The modified suggestion after the first few correct words are typed.

Fig. 12 Interactive translation prediction in Lilt

the suggestions. In fact, the initial suggestion may be rather inadequate. But as soon as the translator starts typing the MT suggestion changes, adapting dynamically to her input. For example, the translator can overwrite the initial suggestion, enforce a certain start of the target sentence (as in our example of "segment No. 252"), and let the system adapt to it (Fig. 12). If the translator likes the suggestion for the next word or phrase she can simply insert it with a single click. TM matches $\geq$ 75% and TB lookup are still available. But in the Lilt environment they are assigned a secondary role being in effect subsumed by the MT system which takes a center stage.

The most advanced and impressive feature of Lilt is *continuous online re-training* or *adaptation* of its MT engine. As soon as the translation of an active segment is confirmed it, of course, goes into the TM; but *the MT system almost immediately learns from it* too. The same happens when a new term is added to the term base; it becomes available for traditional TB lookup but can also be used almost right away in the MT output. This means that new MT suggestions will be based on the translator's earlier input; the system learns them more or less on the fly.

Clearly this model is very different from PEMT. In the interactive and adaptive translation process the translator's biological mind and language technology tools provide two complementary, closely-coupled resources that can leverage each other's strengths in a truly synergistic fashion:

> When the translator actively works with [a neural machine translation] engine to create a translation, they are able to build and learn from each other, the engine offering up a translation the human may not have considered, and the human serving as a moderator, and in so doing, a teacher of the engine (Healy 2018).

This kind of human–computer interaction presents a remarkable case of *distributed predictive processing*. There is ample evidence from cognitive science research that our brains are *precision*- and not *recall*-oriented. We are much better at selecting the best candidate from a list of suggestions than at generating new suggestions from scratch. Search engines, smartphones, and personal assistants are all tuned to this feature of our minds. ITP and adaptive MT make use of it by transferring a good deal of cognitive load associated with creating new, if often half-baked, translation drafts to computers and letting humans do what they are arguably best at.

While this novel type of human–computer interface has not yet been widely adopted in the translation industry, it has considerable potential. And it is ripe for theoretical investigation. The fact that human language processing blends so smoothly and efficiently with linguistic computation in modern CAT tools, and especially in integrated TM-MT systems, raises intriguing questions about what exactly goes on in our heads when we translate, with and without the aid from computers. In the next section I review recent advances in *cognitive translation studies* and discuss potential new avenues of interdisciplinary research into the nature of the translator's extended mind which would require closer collaboration among translation scholars, professional translators, linguists and philosophers of language, cognitive scientists, neuroscientists, and experts in natural language processing.

7 Opening a Window into the Translator's Extended Mind

7.1 A Host of Questions

Recall the "rational reconstruction" of the translation of a sample sentence from the pharmaceutical domain in Sect. 3, beginning with important syntactic decisions followed by span pre-translation and local semantic interpolation. As was shown in Sect. 3, some of these stages can be partially offloaded to a CAT tool equipped with good bilingual resources (TM and TB). I contended that this entire extended system—comprising one or more human translators, the CAT software, and the resources—is capable of performing highly distributed computational tasks (including automatic fuzzy match repair and fragment assembly) leading to the production of the target sentence. I then suggested that it is only the system as a whole that is responsible for representing the meaning of the source sentence in the target language.

Notably, all the computational subtasks outsourced by the translator to a CAT tool rely on the classical "logic-based" symbolic processing,[39] which accounts for their considerable transparency. Indeed the sorts of syntactic and semantic manipulations involved in fuzzy match repair and fragment assembly performed by CAT tools closely mimic what human translators would do (and in fact what they *did* in the not-so-distant past) in the absence of computers, just much more slowly and less consistently.[40] The situation may be interestingly different in state-of-the-art interactive translation prediction and adaptive machine translation (Sect. 6.2) where a human translator is closely coupled with a neural MT system capable of continuous online re-training and adaptation based on the translator's input. Such systems rely on decidedly "sub-symbolic" and increasingly complex deep-learning neural-network architectures whose performance is both notoriously inscrutable and "unreasonably efficient."[41] The extent to which they can be said to *mimic* (rather than *complement*) human linguistic processing is much less clear.

This raises a host of questions about the operation of the translator's extended mind. For example, the human–computer interaction in the CAT process was analyzed above in explicitly semantic terms (Sects. 4 and 5). But is this really justified? And does it apply to all the subtasks involved in the process—such as fuzzy match repair, smart fragment assembly, and span pre-translation—or only to some? Whether automatic fuzzy match repair or fragment assembly offload distinctively semantic processing may be a nontrivial empirical matter. And the same may be true of other CAT subtasks. Do they involve meaning computation? Or are they purely syntactic symbol manipulations? If so, what makes very similar processes involved in computer-*unaided* human translation different? Does translating *numbers* require any kind of conceptual mediation?[42] If so, does automatic translation of numbers and named numeric entities (dates, times, physical and demographic quantities, and so forth) in CAT tools reduce the semantic load in human bilingual processing?

Many of the same questions arise in the context of adaptive MT. But there are also some new questions. For example, does interactive translation prediction have added cognitive value? Or is it merely an ergonomic, time-saving feature of the algorithm? Are neural-network based predictions syntactically, semantically or cognitively

[39] Such is the nature of the language technologies employed in traditional, "non-smart" CAT tools.

[40] The tools, therefore, fully deserve their portion of *epistemic credit* for the shared task: "If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process" (Clark and Chalmers 1998: 8).

[41] The architectures currently used in machine translation range from *long short-term memory*- and *gated recurrent units*-based sequence-to-sequence encoder-decoder models with attention (Bahdanau et al. 2014) to the more recent *transformer* models (Vaswani et al. 2017) which, at the time of writing, continue to break all performance records. The "unreasonable effectiveness" of neural machine translation is a topic that deserves separate scrutiny but lies beyond the scope of the present paper.

[42] As suggested by recent behavioral studies (Duyck and Brysbaert 2008).

different from symbolic-processing based predictions at work in traditional CAT tools? Is continuous online learning and domain adaptation in "smart" CAT systems, such as Lilt (Sect. 6.2) or SCATE (Vandeghinste et al. 2019), a good model of human learning of new terms and phrases "on the fly"? Is the amount of human syntactic parsing significantly reduced in this adaptive environment due to ongoing sentence completion suggestions from the integrated MT engine? Do these translation suggestions go beyond what the human may have considered (Healy 2018)? If so, is this because they are based on sub-symbolic connections that are learned by an artificial neural network but are not clearly grasped, processed, or even recognized as such by the human translator? Does *sublexical* processing[43] positively contribute to this process?

Questions of this sort proliferate. And there remains a more general prior concern stated at the end of the previous section: what happens in the bilingual human brain when translation is carried out, with or without a CAT tool, smart or not? Translation scholars, NLP experts, and neuroscientists acknowledge that we still know very little about it.[44] Some progress, however, has recently been made and more is forthcoming. Below I briefly review the relevant milestones in cognitive translation studies and raise open questions for more focused interdisciplinary research into the nature of the translator's extended mind.

7.2 Cognitive Translation Studies: Milestones and Prospects

7.2.1 Corpus-Based Studies

Cognitive translation studies (CTS) aim to reconstruct the complex network of processes involved in translation from a vast and diverse array of data. Accordingly, most of this data is associated with translation *processes*, not products. However, in some cases valuable data about processes can be mined from translation products. *Corpus-based* CTS attempt to identify typical target text production patterns and the associated planning and cognitive efforts by studying large parallel corpora of human translations which sometimes include revision history and annotations.[45]

Despite the limitations of this approach[46] it has the advantage of relying on easily available huge volumes of parallel corpora. This makes it potentially useful for comparative analyses of MT- and TM-driven translation processes both of which are, after all, corpus-based.[47] For example, the outputs of raw MT, PEMT, and

[43] Such as that implemented in *bytepair encoding* (Sennrich et al. 2016).

[44] "The neurological mechanisms involved in translating and interpreting are one of the chief known unknowns in translation studies" (Tymoczko 2012: 83). "We do not know much about the cognitive processes involved in the translation task" (Poibeau 2017: 22). "In the vast edifice of [translation and interpreting studies], the neurocognitive room so far amounts to little more than a dark, forlorn attic" (García 2019: 1).

[45] See Alves and Vale (2017).

[46] One limitation has to do with the "inverse problem" known in philosophy of science as *underdetermination*: a given pattern in a translation product may have been generated by different cognitive processes.

[47] Machine translation has been corpus-based since 1990s. And good translation memories are in essence bilingual corpora often annotated with useful metadata.

adaptive MT can be rated against the output of a CAT process using automatic scoring metrics[48] currently employed in the MT industry and research. More specialized methods can be developed for probing lexical consistency in these outputs, detecting major syntactic movements enforced by human translators on the MT and TM suggestions, and even checking the sublexical morphological coherence of the latter to see how much effort, and what kind of effort, the human translator had to expend to correct various morphosyntactic issues in pre-translation.

The majority of approaches in CTS are, however, very explicitly *process*-based. I consider them next.

7.2.2 "Rationalist" Approaches and Think-Aloud Protocols

Earlier CTS[49] were characterized by the popularity of highly theoretical models of the translation process inspired by formal linguistics, most and foremost by generative and transformational grammars. The resulting picture, curiously similar to the famous *Vauquois triangle* (Vauquois 1968) of machine translation, divides the human translation process into the *analysis* stage (whereby the deep structure and, ideally, the logical form of the source sentence are laid bare), the *transfer* stage (whereby the deep source language structure is transferred to the deep target language structure), and the *generation* stage (whereby the target deep structure is transformed into an acceptable surface form). The amount of transfer needed to bridge the difference between the source and target languages decreases with the increase in the amount of analysis done at the previous stage and culminates asymptotically in a fully deverbalized "interlingual" representation.

There is certainly something right about this picture. But its skeletal simplicity precludes it from making direct contact with the diversity and complexity of the empirical data abundantly present in real-life translation scenarios. To tap this data, translation scholars started in the late 1980s to adapt *thinking-aloud protocols* (TAP) earlier developed in cognitive psychology, to study the translation process by asking the translator to verbalize her current thoughts (before they disappear from short-term memory), recording such concurrent verbalizations, and then analyzing them with a variety of coding and classification techniques. The main achievements of this method lie in viewing translation as a series of *problem-solving* steps, which throws light on various strategies deployed by translators to identify translation units, perform bilingual lexical search and cross-linguistic semantic analysis, monitor target language production, and edit completed candidate translations.[50] The limitations of TAP have to do with the highly subjective and self-selective nature of the verbalization data and the vast amount of irrelevant details present in it, which creates problems for generalization. It also raises a more philosophical question: does introspection necessarily provide a good window into the translator's mind?

48 Such as BLEU, TER or HTER.

49 See e.g. Nida (1964).

50 For a review and critical discussion of TAP-based CTS, see Bernardini (2001).

Methodologically speaking, rationalist[51] and TAP-based approaches seem to pull in opposite directions: the former tend to impose top-down, almost normative constraints on what a translation process should look like; the latter embrace a bottom-up empirically-grounded methodology and sensitivity to details. Perhaps some combination of these methods can strike a good balance. I submit that the reconstruction of the CAT process in Sects. 3 and 4 and subsequent philosophical reflection on it in Sect. 5, informed as they are by general semantic considerations as well as the first-person perspective of a translator familiar with CAT tools, combine both approaches. Be that as it may, any conclusions drawn from these two sources are tentative and must be confirmed or disconfirmed by more objective empirical data, to which I now turn.

7.2.3 Keylogging and Eye-Tracking

Keylogging programs record and time-stamp keyboard strokes and mouse movements of a computer user as she performs a given task. Shortly after their introduction in the early 1990s such programs were adapted for CTS,[52] which resulted in a steady flow of new objective data. The use of keylogging in cognitive studies of any kind is predicated on the assumption that typing behavior is indicative of concurrent mental processing. Modern keylogging techniques[53] extended with pre-processing algorithms allow researchers to filter out noise and identify composite operations such as character and word insertions, deletions, and replacements which combine multiple keystrokes and mouse movements. Keylogging data led to statistically robust inferences about the distribution of cognitive effort and attention across translation stages, such as initial skimming, drafting and revision, and various dependencies such as that between the difficulty of the source document and the optimal length of translation units which the user chooses to handle one at a time. Recent application of keylogging to interactive and adaptive machine translation yielded useful measurements of the temporal and technical effort (processing time, manual insertions and deletions, mouse-driven navigation) involved in ITP versus PEMT.[54]

While target text production can be studied with keylogging, source text comprehension can be monitored with *eye-tracking*—another powerful method imported into CTS from usability research and other applications.[55] According to the widely accepted "eye-mind assumption,"[56] data such as gaze fixation, saccade and regression tracing, and pupil dilation can be used to make hypotheses about the intensity, duration, and timing of various cognitive sub-processes involved in translation. This technique has been applied with varying degree of success to quantify cognitive effort in tasks as different as processing fuzzy matches in CAT tools, where the

[51] The term is García's (2019: Sect. 1.2.1).

[52] See Jakobsen (1999).

[53] Such as those implemented in Translog-II (Carl 2012) and Casmacat (Sanchis-Trilles et al. 2014).

[54] Knowles et al. (2019).

[55] For an overview, see Göpferich et al. (2008).

[56] "There is no appreciable lag between what is being fixated and what is being processed" (Just and Carpenter 1980: 331).

processing time was plotted against fuzzy match rates based solely on the pupil dilation data, and evaluating the quality of MT output.[57]

It is not so clear that eye-tracking by itself can reveal the most intimate and linguistically important distinctions, such as that between syntactic and semantic bilingual processing or between lexical access and sentence or phrase comprehension.[58] But it is obvious that eye-tracking can usefully contribute to this effort, especially when it is deployed in combination with other methods. In fact, eye-tracking and keylogging are increasingly perceived by translation scholars as two sides of the same cognitive coin.[59] Moreover, source language comprehension and target language production often overlap in time, especially in a professional setting. This is obviously true of simultaneous interpreting. However, recent studies employing both eye-tracking and keylogging make it clear that these two cognitive phenomena "cannot easily be separated as two distinct activities" even in written translation (Dragsted 2010: 43).[60] New conceptual frameworks have been developed to describe the difference between *sequential* and *integrated coordination* of reading and writing[61] processes in translation (*ibid*.), *horizontal* versus *vertical* translation, *serial* versus *parallel* translation, and between *alternating* and *divided attention*.[62] The underlying idea is that translation may begin before the source sentence is fully understood.

While these findings emerge from studies of computer-unaided translation they seem to be relevant to the CAT process where, as argued in Sect. 5, the source sentence may *never* be fully understood by the individual translator. Instead, understanding is *divided* among the synergistic components of a system comprising one or more human translators, their common CAT tool, and sometimes an MT module. Tracking the translator's visual attention and typing activity in such circumstances based on carefully designed metrics could substantially contribute to better appreciation of the role of the closely coupled but still physically separate elements of his extended mind.

As an example, recall *span pre-translation*, which does so much heavy lifting in the CAT process (Sect. 3.2), with its "linguistically unaware" strings (such as *cell types, including*) cutting across phrase boundaries, and the cognitive load associated with mentally aligning their source and target counterparts. Just how burdensome is such alignment? It probably depends on whether a given string of this sort can be stored in the mind and manipulated as a semiautonomous unit. In their everyday work translators quickly learn various heuristic shortcuts which allow them to minimize the cognitive cost of dealing with complex sentences. A prominent NLP expert

[57] See, respectively, O'Brien (2008) and Doherty et al. (2010).

[58] For some doubts on this score, see García (2019: 15).

[59] See Carl and Kay (2012) and Schaeffer et al. (2019).

[60] The similarity between simultaneous interpreting and "parallel" or "vertical" professional translation is reflected in the terms introduced by CTS scholars to describe the time lag between the source text input and the start of its interpreting in the former (*ear-voice span*) and the time lag between a fixation on a source word and the first keystroke associated with its translation (see Dragsted 2010).

[61] i.e., *typing*. Notably, many professional translators are *touch typists*; they can type into the target window while being fixated on the source window.

[62] See Carl and Kay (2012: 954).

notes, in a similar connection, that "even high-level syntactic structures can correspond to regular patterns… or 'prefabs'":

> In this framework, syntax is not as prominent as in traditional approaches; the sentence is seen as an assemblage of "prefab units," or, put differently, an assemblage of complex sequences stored as such in the brain. The analysis is thus simpler, since, if this hypotheses is correct, the brain does not really have to take into account each individual word but has direct access to higher-level units, reducing both the overall ambiguity and the complexity of the sentence-understanding process (Poibeau 2017: 23).

The extent to which syntactically ill-formed pre-translated spans can function as such "prefabs" might be determined by comparing the cognitive load involved in their processing and manipulation with the corresponding variables for genuine grammatical phrases. This could be done with a combination of eye-tracking and keylogging methods similar to those that allowed the authors of a recent influential study to quantify the processing of various "translation units"—bilingual "units of cognitive activity"—and to dissociate them from "alignment units" that reflect semantic correspondences in the finished translation products.[63]

To be sure, evaluating fine-grained parameters of bilingual processing in the CAT or adaptive MT environment is not easy, even with a combination of eye-tracking and keylogging. In addition to the identification of the relevant variables and design of ecologically valid experimental conditions, CAT-customized bilingual testing resources must be carefully prepared. Another challenge has to do with integrating existing keylogging and eye-tracking programs with the interfaces of CAT tools and/or adaptive MT systems that are actually used by translators (as opposed to those that are specifically designed for research purposes), to ensure ecological validity. Some promising attempts in this direction, however, have been made.[64] Alternatively, one could try to recreate CAT-specific processes such as fuzzy match repair and fragment assembly in an artificial testing environment.[65] In all cases, this would require a joint effort of linguists, psychologists, NLP experts, software engineers, translation scholars and, of course, professional translators.

7.2.4 Psycholinguistic Studies of Bilingualism and the Neuroscience of Translation

The foregoing brief review of CTS has focused on several interrelated theoretical and behavioral lines of inquiry that currently appear to have the greatest potential for probing the details of the computer-enhanced human translation process. One might think of these core methods as occupying a somewhat gerrymandered area on the map of cognitive studies more generally. This area, however, has vague boundaries where it blends into and is continuously informed by several much better-established research industries.

[63] Carl and Kay (2012).

[64] For example, Teixeira and O'Brien (2017) report the results of their study of the cognitive ergonomic aspects of memoQ (see Sect. 2 above) using keylogging, eye-tracking, and screen recording.

[65] Such as Translog II (Carl 2012) equipped with reliable gaze-to-word mapping software.

On the one hand, there are *psycholinguistic studies of bilingualism* with a long tradition of experimental methods for exploring various aspects of bilingual processes—morphological, syntactic, semantic, and pragmatic—based on carefully designed behavioral tasks in which participants are asked to evaluate equivalence relations between different types of words (verbs vs. nouns, abstract vs. concrete, cognate vs. non-cognate, etc.), phrases or sentences in their two languages (L_1 and L_2), or to perform short translations in both directions, *forward* (FT: $L_1 \rightarrow L_2$) and *backward* (BT: $L_2 \rightarrow L_1$). Participants may be suitably primed, and their response time and accuracy logged for analysis. Studies of this sort, conducted on healthy as well as aphasic subjects, have revealed important dissociations between FT and BT processing in bilinguals, between syntactic and semantic operations during translation, and between the cognitive costs of translating various parts of speech, phrase types, and thematically different constructions.[66] Along with related behavioral metrics for evaluating executive and other non-verbal cognitive functions (such as working memory and attention management, inhibitory control, and mental set shifting), psycholinguistic approaches can be extended to experimental situations more closely reproducing the CAT process. In fact, some of this has already been done indirectly in the context of eye-tracking and keylogging experiments mentioned above—a clear indication that the boundary between the two methods is indeed vague.

On the other hand, all higher cognitive processes without exception take place in the brain. In an obvious sense, the *head* of the translator is the natural object of study for anyone who wants to understand what goes on in the translator's mind. It must be acknowledged that translation per se has never been a priority for neurolinguists. There are several reasons for this, including (i) the overabundance of "bigger" open questions on the agenda of neurolinguistics, (ii) the extreme rarity of the relevant clinical cases (such as professional translators developing the relevant types of aphasia), and (iii) the formidable obstacles to recreating realistic translation environment (let alone CAT environment!) in ERP and fMRI settings.[67] Accordingly, most of the neuroscientific data relevant to translation comes from the broader studies of the bilingual brain.

But a growing number of players on both sides of the "two-culture divide" are emphasizing the importance of more focused investigation of the neural substrates of the translation processes.[68] In a recent review of the neurocognition of translation and interpreting, Adolfo García[69] notes remarkable achievements in the identification of partially dissociable brain networks for different translation directions in bilinguals (BT versus FT), processing levels (syntactic versus semantic), and translation unit types (sentences versus words),[70] which can be aligned with the similar

---

[66] See García (2015).

[67] In contrast, it is much easier to recreate realistic *simultaneous interpreting* scenarios for the purpose of ERP and fMRI studies.

[68] See Tymoczko (2012) and García (2019).

[69] Who exemplifies both cultures: initially trained as a scientific translator, he is both a neurolinguist and a translation scholar.

[70] According to some recent data, FT tends to generate greater modulations in Broca's area and the putamen than BT; translation of sentences and of action verbs elicits more activity in the frontostriatal regions, and translation of other words (and especially concrete nouns) in the temporal-parietal regions. Importantly, some of these distinctions can be observed in the absence of explicit behavioral effects (García 2019: 207).

evidence from behavioral studies (García 2019: 206–207). Along with the richer data from neurocognitive studies of interpreting processes, these empirical results have led to detailed neuroanatomical models of bilingual processing which are open to confirmation or refutation. This helps to balance the sometimes too "revolutionary" trends in traditional theoretical translation studies with the *normal* (in Kuhnian terms), continuous and steady development of neurocognitive research of bilingualism.[71] Last but not least, cognitive translation studies could greatly benefit from incorporating neuro-computational models and insights from the recent advances in computational neurolinguistics.[72]

Some of these developments could certainly shed more light on the operation of the translator's extended mind. Philosophers of science and of language could usefully contribute to this interdisciplinary effort by applying their distinctive conceptual toolkits to study the epistemic, methodological, and semantic dimensions of computer-enhanced human translation. Indeed, the latter could become a new testing ground for theoretical models of human language processing. For example, there may be an interesting question whether the eye-key span in the CAT process and the ear-voice span in simultaneous interpreting implicate the same amodal neural substrates or different modality-specific circuits. It may also be of interest to explore the extent to which translation predictions generated by sub-symbolic processing in adaptive neural MT are semantically transparent (or not) for humans. There are many other intriguing questions to raise. But we must leave the matter here.

8 Conclusion

The highly distributed nature of the human–computer interaction in computer-assisted translation can be explored from various angles, from theoretical to decidedly empirical. In our case study of the CAT process we focused on its distinctly semantic aspects demonstrating the translator's extended mind at work in the context of real-life scenarios. Then we reviewed the latest developments in adaptive machine translation which aim to synergize the complementary computational powers of the biological mind and artificial neural networks. Recent advances in cognitive translation studies, from theoretical modeling to experimental paradigms, demonstrate the potential for probing the details of these processes. This opens promising new avenues of interdisciplinary research which would require closer collaboration among professional translators, translation scholars, linguists and philosophers, cognitive scientists and neuroscientists.

Acknowledgements Work on this paper was completed during the author's "study in a second discipline" (linguistics) in AY 2019–20. The author thanks UGA's Office of the Senior Vice President for Academic

[71] Cf. García (2019: 40, 209).

[72] See, in particular, Buchweitz et al. (2012), Pereira et al. (2018), and Abnar et al. (2018).

Affairs and Provost, the Franklin College of Arts and Sciences, and the Departments of Philosophy and Linguistics for their support. The author thanks Galia Williams and Konstantin Lakshin for discussions of some issues raised in the paper, and the reviewers for the very helpful comments.

Compliance with Ethical Standards

Conflict of interest The author declares no conflict of interest.

References

Abnar, S., Ahmed, R., Mijnheer, M., & Zuidema, W. (2018). Experiential, distributional and dependency-based word embeddings have complementary roles in decoding brain activity. *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, 57–66. https://doi.org/10.18653/v1/W18-0107.

Alves, F., & Vale, D. C. (2017). On drafting and revision in translation: A corpus linguistics oriented analysis of translation process data. In *Annotation, exploitation and evaluation of parallel corpora* (pp. 89–110). Language Science Press. https://doi.org/10.5281/ZENODO.283500.

Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. Retrieved from arXiv: https://arxiv.org/abs/1409.0473.

Bar-Hillel, Y. (1960). The present status of automatic translation of languages. *Advances in Computers, 1,* 91–163.

Barrachina, S., Bender, O., Casacuberta, F., Civera, J., & Vilar, J.-M. (2009). Statistical approaches to computer-assisted translation. *Computational Linguistics, 35*(1), 3–28. https://doi.org/10.1162/coli.2008.07-055-R2-06-29.

Bernardini, S. (2001). Think-aloud protocols in translation research: Achievements, limits, future prospects. *Target, 13*(2), 241–263. https://doi.org/10.1075/target.13.2.03ber.

Buchweitz, A., Shinkareva, S. V., Mason, R. A., Mitchell, T. M., & Just, M. A. (2012). Identifying bilingual semantic neural representations across languages. *Brain and Language, 120*(3), 282–289. https://doi.org/10.1016/j.bandl.2011.09.003.

Carl, M. (2012). Translog-II: A program for recording user activity data for empirical translation process research. *LREC, 12,* 4108–4112.

Carl, M., & Kay, M. (2012). Gazing and typing activities during translation: A comparative study of translation units of professional and student translators. *Meta, 56*(4), 952–975. https://doi.org/10.7202/1011262ar.

Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. Oxford: Oxford University Press.

Clark, A. (2014). *Mindware: An introduction to the philosophy of cognitive science* (2nd ed.). Oxford: Oxford University Press.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis, 58*(1), 7–19.

Daems, J., & Macken, L. (2019). Interactive adaptive SMT versus interactive adaptive NMT: A user experience evaluation. *Machine Translation, 33*(1–2), 117–134. https://doi.org/10.1007/s10590-019-09230-z.

Doherty, S., O'Brien, S., & Carl, M. (2010). Eye tracking as an MT evaluation technique. *Machine Translation, 24*(1), 1–13. https://doi.org/10.1007/s10590-010-9070-9.

Dragsted, B. (2010). Coordination of reading and writing processes in translation: An eye on uncharted territory. In G. M. Shreve & E. Angelone (Eds.), *Translation and cognition* (pp. 41–62). Amsterdam: John Benjamins Publishing Company. https://doi.org/10.1075/ata.xv.04dra.

Duyck, W., & Brysbaert, M. (2008). Semantic access in number word translation: The role of crosslingual lexical similarity. *Experimental Psychology, 55*(2), 102–112. https://doi.org/10.1027/1618-3169.55.2.102.

Evans, G. (1973). The causal theory of names. *Aristotelian Society Supplementary, 47*(1), 187–208. https://doi.org/10.1093/aristoteliansupp/47.1.187.

García, A. M. (2015). Psycholinguistic explorations of lexical translation equivalents: Thirty years of research and their implications for cognitive translatology. *Translation Spaces, 4*(1), 9–28. https://doi.org/10.1075/ts.4.1.01gar.

García, A. M. (2019). *The neurocognition of translation and interpreting*. Amsterdam: John Benjamins Publishing Company.

Göpferich, S., Jakobsen, A. L., & Mees, I. M. (Eds.). (2008). *Looking at eyes: Eye-tracking studies of reading and translation processing*. Samfundslitteratur.

Green, S., Chuang, J., Heer, J., & Manning, C. D. (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 - UIST'14*, (pp. 177–187). https://doi.org/10.1145/2642918.2647408

Green, S., Heer, J., & Manning, C. D. (2015). Natural language translation at the intersection of AI and HCI. *Communications of the ACM, 58*(9), 46–53. https://doi.org/10.1145/2767151.

Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C, Huang, X., Junczys-Dowmunt, M., Lewis, W., Li, M., Liu, S., Liu, T. Y., Luo, R., Menezes, A., Qin, T., Seide, F., Tan, X., Tian, F., Wu, L., Wu, S., Xia, Y., Zhang, D., Zhang, M., & Zhou, M. (2018). Achieving human parity on automatic Chinese to English news translation. ArXiv: 1803.05567 [Cs]. http://arxiv.org/abs/1803.05567.

Healy, C. (2018). Here's Why Neural Machine Translation is a Huge Leap Forward. Retrieved June 23, 2020, from https://labs.lilt.com/neural-machine-translation-huge-leap-forward.

Hohwy, J. (2013). *The predictive mind*. Oxford: Oxford University Press.

Hutchins, W. John. (1986). *Machine translation: Past, present, future*. Ellis Horwood: Halsted Press.

Hutchins, J. (1998). The origins of the translator's workstation. *Machine Translation, 13*(4), 287–307.

Hutchins, W. J., & Somers, H. L. (1992). *An introduction to machine translation*. New York: Academic Press.

Jakobsen, A. L. (1999). Logging target text production with Translog. *Copenhagen Studies in Language, 24,* 9–20.

Jurafsky, D., & Martin, J. H. (2008). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition* (2nd ed.). Upper Saddle River: Prentice Hall.

Just, M. A., & Carpenter, P. A. (1980). A theory of reading: From eye fixations to comprehension. *Psychological Review, 87*(4), 329–354. https://doi.org/10.1037/0033-295X.87.4.329.

Kay, M. (1980). *The proper place of men and machines in language translation* [Research Report CSL-80-11]. CA: Xerox Palo Alto Research Center.

Kay, M. (2017). *Translation: Linguistic and philosophical perspectives*. Stanford: Center for the Study of Language and Information.

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science, 18*(4), 513–549. https://doi.org/10.1207/s15516709cog1804_1.

Knowles, R., Sanchez-Torron, M., & Koehn, P. (2019). A user study of neural interactive translation prediction. *Machine Translation, 33*(1–2), 135–154. https://doi.org/10.1007/s10590-019-09235-8.

Koehn, P. (2010). *Statistical machine translation*. Cambridge: Cambridge University Press.

Koehn, P. (2020). *Neural machine translation*. Cambridge: Cambridge University Press.

Nida, E. A. (1964). *Toward a science of translating: With special reference to principles and procedures involved in bible translating*. Netherlands: Brill.

O'Brien, S. (2008). Processing fuzzy matches in translation memory tools: An eye-tracking analysis. In Göpferich et al. (pp. 79–102).

Pereira, F., Lou, B., Pritchett, B., Ritter, S., Gershman, S. J., Kanwisher, N., et al. (2018). Toward a universal decoder of linguistic meaning from brain activation. *Nature Communications, 9*(1), 963. https://doi.org/10.1038/s41467-018-03068-4.

Poibeau, T. (2017). *Machine translation*. Cambridge: The MIT Press.

Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science, 7,* 131–193.

Sánchez-Gijón, P., Moorkens, J., & Way, A. (2019). Post-editing neural machine translation versus translation memory segments. *Machine Translation, 33*(1–2), 31–59. https://doi.org/10.1007/s10590-019-09232-x.

Sanchis-Trilles, G., Alabau, V., Buck, C., Carl, M., & Vidal, E. (2014). Interactive translation prediction versus conventional post-editing in practice: A study with the CasMaCat workbench. *Machine Translation, 28*(3), 217–235. https://doi.org/10.1007/s10590-014-9157-9.

Schaeffer, M., Nitzke, J., Tardel, A., Oster, K., Gutermuth, S., & Hansen-Schirra, S. (2019). Eye-tracking revision processes of translation students and professional translators. *Perspectives, 27*(4), 589–603. https://doi.org/10.1080/0907676X.2019.1597138.

Searle, J. (1980). Minds, Brains, and Programs. *Behavioral and Brain Sciences, 1,* 417–424.
Sennrich, R., Haddow, B., & Birch, A. (2016). Neural machine translation of rare words with subword units. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (pp. 1715–1725). https://doi.org/10.18653/v1/P16-1162.
Sin-wai, C. (2017). *The future of translation technology: Towards a world without Babel*. London and New York: Routledge.
Somers, H. (2003). Translation memory systems. In H. Somers (Ed.), *Computers and translation: A translator's guide* (pp. 31–47). Amsterdam: John Benjamins Publishing Company. https://doi.org/10.1075/btl.35.06som.
Teixeira, C. S. C., & O'Brien, S. (2017). Investigating the cognitive ergonomic aspects of translation tools in a workplace setting. *Translation Spaces, 6*(1), 79–103. https://doi.org/10.1075/ts.6.1.05tei.
Tiedemann, J. (2012). Parallel Data, Tools and Interfaces in OPUS. In K. Choukri, T. Declerck, M. U. Doğan, & S. Piperidis (Eds.), *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*. Istanbul, Turkey: European Language Resources Association (ELRA), pp. 2214–2218.
Tymoczko, M. (2012). The neuroscience of translation. *Target, 24*(1), 83–102. https://doi.org/10.1075/target.24.1.06tym.
Vanallemeersch, T., & Vandeghinste, V. (2015). Assessing linguistically aware fuzzy matching in translation memories. *Proceedings of the 18th Annual Conference of the European Association for Machine Translation*, 153–160. https://www.aclweb.org/anthology/W15-4920.
Vandeghinste, V., Coppers, S., Van den Bergh, J., Bulté, B., Rigouts Terryn, A., Lefever, E., van der Lek, I., Coninx, K., & Steurs, F. (2017). The SCATE prototype: A smart computer-aided translation environment. *Proceedings of the 39th Conference on Translating and the Computer*, AsLing, pp., 104–113.
Vandeghinste, V., Vanallemeersch, T., Augustinus, L., Bulté, B., & Luyten, K. (2019). Improving the translation environment for professional translators. *Informatics, 6*(2), 24. https://doi.org/10.3390/informatics6020024.
Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30* (pp. 5998–6008). Curran Associates, Inc. http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf.
Vauquois, B. (1968). A survey of formal grammars and algorithms for recognition and transformation in mechanical translation. In A. J. H. Morrel (Ed.), *IFIP Congress (2)* (pp. 1114–1122). http://dblp.uni-trier.de/db/conf/ifip/ifip1968-2.html#Vauquois68.
Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Dean, J. (2016). Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. https://arxiv.org/abs/1609.08144.
Wuebker, J., Green, S., DeNero, J., Hasan, S., & Luong, M.-T. (2016). Models and inference for prefix-constrained machine translation. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (pp. 66–75). https://doi.org/10.18653/v1/P16-1007.