Moving Towards Personalising Translation Technology

Sharon O'Brien is Associate Professor of Translation Studies at the School of Applied Language and Intercultural Studies, Dublin City University. She is a funded investigator at the Science Foundation Ireland-funded research centre, ADAPT.

Owen Conlan is Associate Professor in the School of Computer Science and Statistics of Trinity College Dublin. His primary research interests are in personalisation and visualisation. Owen leads the Personalising the User Experience research theme as part of the Science Foundation Ireland-funded ADAPT research centre.

Abstract

Technology has had an important impact on the work of translators and represents a shift in the boundaries of translation work over time. Improvements in machine translation have brought about further boundary shifts in some translation work and are likely to continue having an impact. Yet translators sometimes feel frustrated with the tools they use. This chapter looks to the field of personalisation in information technology and proposes that personalising translation technology may be a way of improving translator-computer interaction. Personalisation of translation technology is considered from the perspectives of context, user modelling, trust, motivation and well-being.

Keywords

Computer-Aided Translation, Machine Translation, Personalisation, Adaptation, User Modelling

Acknowledgement

This research is funded by Science Foundation Ireland. The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

Moving Towards Personalising Translation Technology

Sharon O'Brien,
School of Applied Language
and Intercultural Studies,
ADAPT Research Centre,
Dublin City University

Owen Conlan,
School of Computer Science,
and Statistics,
ADAPT Research Centre,
Trinity College Dublin

Introduction

The translation profession has seen many changes since the 1990s, notably since the introduction of Translation Memory (TM) tools. It is generally accepted that the 'technological turn' in translation (O'Hagan 2013) has had an impact on translators, on the translation process and product, as well as on the academic discipline of Translation Studies. The more recent successes of data-driven statistical machine translation (SMT) and, even more recently still, neural machine translation (NMT) signal yet another change for some sectors of the translation profession. Generally speaking, technologisation is driven by the basic assumption that things will be better as a result. Few would argue that translation memory tools, for instance, have not brought about advantages. Yet recent research shows that there is still disgruntlement among users of TM technology. And, despite significant advances in Machine Translation (MT), the voices within the translation profession singing praises for that technology are few.

This chapter first examines these shifts in the translation technology landscape and argues that the traditional boundaries between TM and MT technology are becoming blurred, so much so that it is difficult now to treat them as separate. This is relevant for translation practice, research and teaching. Where a clear divide existed previously, that is now disappearing. It considers the implications of increasing technologisation, especially in the context of some continued dissatisfaction among translators with their tools. The paper then turns its attention to the concepts of personalisation and adaptation to explore whether and how these notions could be applied, to the benefit of users, to translation technology. It is emphasised here that personalisation and adaptation are not trivial techniques and so we do not suggest trivial tweaks to existing technologies, but fundamental changes that could be explored to see if personalising translation technology might eventually lead to a better symbiosis between the translator and her tools. As personalisation and adaptation have not, to the best of our knowledge, been considered at any length in the context of translation, the ideas presented here are necessarily exploratory.

Major Boundary Shifts

Translation Memory

The introduction of TM tools in the early 1990s represented a considerable shift in translator work practices and had an important effect on the translation process. Prior to its introduction, in high volume, repetitive translation environments, the word processing 'compare' feature was used to identify changes in the source text and translations were then cut and pasted into the updated source, before translation could continue. This tedious and error-prone process was eliminated by TM tools. Expectations for translator productivity consequently increased as did expectations for quality, in particular for the dimension of consistency. In the TM environment, translators now had to deal with processing two source texts (ST), the 'true' ST and a match from the TM database. The implications of this were that they not only had to engage with the usual translation sub-processes, but also cross-language comparison and evaluation, acceptability decision making and editing.

In large-scale translation contexts, translation became more 'collaborative', or derivative, since translations from others were stored in a shared database and reused. Furthermore, TM tools meant that translators did not necessarily have to be experts in multiple file formats and applications such as Framemaker, SGML, or HTML. That is, they now translated within the TM environment, which filtered these file formats. There were economic implications too as TM tools forced a downward pressure on cost-per-word based on exact and fuzzy match volumes. It is also argued that TM tools forced translators to focus more on segments than on text and that this impacted on the translated product (Dragsted 2005; Mellinger and Shreve 2016).

Thus it can be argued that TM technology impacted on the profession, the product and the process of translation. TM tools are now firmly embedded in many sectors of the translation profession. In a recent survey by SDL (SDL, Online) one of the leading Computer-Aided Translation (CAT) tool developers, 83% of respondents reported using "translation productivity" software. The survey had 2,784 responses from across 115 countries. Over half of the respondents had at least 5 to more than 10 years of experience with these tools.

We could assume that the major shift that occurred with the introduction of TM is long past, yet more than twenty years on translators still sometimes report that they are not completely satisfied with their TM tools, in particular they mention complexity of the user interface and forced segmentation as being problematic (O'Brien et al. (forthcoming), Moorkens and O'Brien 2016, LeBlanc 2013).

Statistical Machine Translation

The more recent increase in MT usage can be compared to this seismic shift represented by the introduction of TM. Rules Based MT trundled on for many years in the background, being used in only some organisations (e.g. PAHO – the Pan-American Health Organisation – Vasconcellos 1985, Vasconcellos and León 1985). The introduction of data-driven MT, or SMT, on the other hand, resulted in a considerable uptake in the specialised translation market, though this has been less accelerated than for TM, and the impact has been limited to certain domains and language pairs (e.g. IT, legislation, TAUS 2014). The impact of advances in MT technology could be considerably greater in the longer term. Some translation scholars are asking if all translators will become post-editors (Pym 2013) and suggest that translation technologies, including MT, "are altering the very nature of the translator's cognitive activity" (Pym 2011: 1). There will be implications for translator training programmes and for models of translation competence where one additional competence will be to learn to trust the data (Pym 2013).

Neural Network MT?

At the time of writing, a new shift in boundaries for translation technology has already become obvious. SMT has certainly brought about advances in the quality that could be produced by MT systems, but the general consensus is that a quality ceiling had been reached. At the same time, advances in artificial intelligence (AI), in particular in the domain of neural networks, have occurred. More and more (translation) data has become available and the high performance computing requirements for processing neural networks is becoming more of a reality. AI researchers have consequently turned their attention to whether MT quality could be improved by using a neural network, rather than a phrase-based statistical approach.

NNMT stands for Neural Network Machine Translation. NNMT, or just NMT, is based on the concept of deep learning and neural nets. Deep learning is linked with the concept of machine learning, where a computer automatically 'learns' from data. For MT, this means that the system can create a predictive model to translate new source material based on 'knowledge' it has gleaned from translation memory or other natural language data. In contrast, SMT systems work on the word and phrase level, which means that some phrases can be fluently and accurately translated within a sentence, while other parts of the sentence can sound disfluent. The main promise offered by NMT over SMT is that greater context can be taken into account and that this can lead to greater accuracy and fluency when compared with SMT. (For an accessible comparison of NMT and SMT, see Forcada 2017 and the blog post by Vashee 2016).

Although NMT appears promising, it still has limitations at the time of writing. For example, it is reportedly slower than SMT because it is computationally more complex (days vs. weeks of time for training the engine). Also, increases in quality over SMT engines have been reported primarily in the currency of automatic evaluation metrics, such as BLEU scores (see Koehn 2010 for a full description of BLEU and similar metrics), which still do not have acceptance as a valid quality metric among translation scholars or the translation profession. The impact of NMT is unknown at this point, but it is likely to contribute further to the changing translation technology landscape.

Blurring Boundaries

With ever-increasing usage of MT, we are witnessing a blurring of traditional boundaries between the two technologies of TM and MT. Many TM tools now also include MT suggestions as an option. It is expected that eventually the technology will be such that automated confidence estimation (a quality estimate generated by the MT system) will be used to triage MT by comparing with available TM matches and presenting the best suggestion to the translator, which might be a Fuzzy Match from the TM database or a suggestion from the MT system (Dara et al. 2013).

One term that has been used to describe translator interaction with MT is 'Human-in-the-Loop' automated translation, but this places the human translator in a peripheral position in relation to the machine process and reduces her to the role of verifying or fixing the machine errors. A more desirable description from the translator's perspective might be 'Machine-in-the-Loop' translation, where the human agent is served by the machine, and not the other way around.

It is not only the case that MT is now mixed in with TM, but the boundaries between what is 'human translation' and what is 'machine translation' are also becoming blurred. With the application of deep learning techniques, the machine effectively learns from the examples produced by humans (via TMs and parallel translated data). The learning can happen at particular junctures in time moving from intervals of weeks or months, to within the same text editing instance, to within the current segment editing instance. In the first case, the system is retrained using post-edited data at monthly intervals, for example. In the second, the system learns on the fly through what is called 'adaptive' MT and provides suggestions to the translator based on decisions she made earlier in the text editing process. The third interval is known as 'Interactive MT', where the MT proposal for the current segment changes in real time depending on decisions the translator makes. Interactive MT was proposed as a prototype technology many years ago (Foster et al. 1997, Foster et al. 2002) and has lately become a reality through technologies such as Lilt,which has been hailed by some in the translation profession as a game changer (Zetsche 2016). With adaptive and interactive MT, the dividing lines between TM and MT are no longer clear.

Implications of shifting and blurring boundaries?

Considering the changing landscape sketched above and assuming that NMT will meet expectations, it is not unrealistic to state that translation technology will only grow in importance in the translation profession. There are concerns that the role of the human translator will be, at a minimum, reduced to that of 'post-editor'. A more optimistic view is presented by Melby (2016), who argues that until we experience the phenomenon known as the "singularity" (Kurzweil 2010), humans will still outperform machines in terms of emotion and agency, and maybe also creativity. (The 'singularity', as explained by Kurzweil, refers to a time in the future when technological change will be so rapid and so advanced that human life as we know it now will be radically and irreversibly changed.)

We know from experience that the introduction of TM technology represented a significant change for translators and that it was not always positive. That the introduction of TM technology was seen and experienced as being forced by the translation industry in a top-down manner was unhelpful. Translators were unsuspecting and ill-prepared; they were not consulted in terms of needs or design of tools and we are still seeing the repercussions today in terms of dissatisfaction with tools and the process, as mentioned in the section on 'Major Boundary Shifts' above As Olohan puts it, systems sometimes fail (or meet with less success) because their development is seen as a technical change process, rather than a *socio*-technical change process (Olohan 2011, our emphasis).

Considering the major shifts we are witnessing due to technological innovation, and to avoid repeating the mistakes of the past, it is surely worth making translators central to the current developments. More collaboration and consultation between technology researchers and developers and their end users is needed. However, we go beyond calling for just consultation of translators regarding the design of the technologies they use and move one step further to call for intelligent personalisation and adaptation of translation technologies. The next section will discuss this in some detail.

Considerations on personalisation and adaptation

What is personalisation and adaptation?

Prior to discussing personalisation and adaptation in the context of translation (see the Section 'Relevance to Translation Technology' below), we first give a brief introduction to those concepts. According to Göker et al. (2002: 4), "[p]ersonalisation is about tailoring products and services to better fit the user". There are several ways of achieving personalisation, and the main ways involve focusing on the user needs, preferences, interests, expertise, workload, tasks etc. "We advocate user context as a means of capturing all these" (ibid). Personalisation has been shown to have strong potential in allowing users to access and understand complex information and processes (Hampson et al. 2014), and in alleviating cognitive overload (Höök 1998).

At the core of all personalisation research is the need for a User Model (Knutov et al. 2009), which uses terms from a Domain Model to indicate a user's relationship to different concepts. The Domain Model describes how concepts are connected to each other defining a semantic relationship between them (De Bra et al. 1999, Conlan et al. 2002, Brusilovsky 2008). The User Model contains characteristics of individual users such as goals, knowledge, background and preferences modelled in terms of the Domain Model concepts. For example, in the case of translation, the User Model might contain information on the degree of expertise or experience a translator has in translating a specialised domain (background). Information about environment and location is often added to the User Model to handle mobility and the use of various devices (Joerding 1999, Garlatti et al. 1999, Billsus et al. 2000).

Discussion of personalisation is often accompanied by the term *adaptation*, which necessarily raises a question about the differences between these two concepts. García-Barrios et al. (2005: 122) discuss the conceptual differences between personalisation and adaptation and argue that the concepts are interdependent: "*personalising* is the same as *adapting towards a specific user*" and, therefore, "personalisation systems represent a specific subtype of general adaptation systems" (emphasis in original).

Background and Contexts

Early discussions on adaptive user interfaces stem from the 1980s and initially focused on needs, preferences and expertise but then also merged with work on user modelling. IT systems that model users in this way are sometimes referred to as 'context-aware applications' and 'affective user interfaces'. According to Göker and Myrhaug (2002: 5), "[a] context can be defined as a description of aspects of a situation." Personalisation can be implemented in a range of contexts, including, for instance, online e-commerce (e.g. Karat et al. 2004) and e-learning (Green et al. 2005, Conlan et al. 2002).

Green et al. (2005) argue for the reversal of logic in education systems to make the system conform to the learner, rather than the learner to the system. This is, according to them, “the essence of personalisation”. Commenting on the same domain, Oulasvirta and Blom highlight that individuation of learning materials has been shown to increase “not only motivation, but depth of engagement, amount learned, perceived competence, and levels of aspiration” (2008: 3).

Motivation for personalisation

Oulasvirta and Blom (2008: 13) tell us that “there is no special need for personalisation, rather there are context-independent basic needs that are idiosyncratically manifested as motivations related to the use of a product’s features.” Using Deci and Ryan’s Theory of Self-Determination (Deci and Ryan 1985, 2000), Oulasvirta and Blom (2008: 14) claim that “there is a link between well-being and personalisation”. As mentioned above, they also show better learning experience through personalisation in the field of e-learning and list the following as motivational factors for personalisation: autonomy, competence, and relatedness. Autonomy is about freedom and “unpressured willingness” to engage in an activity (ibid: 5). It can be affected by surveillance, evaluation and deadlines. Competence is seen by them as “a psychological need that provides an inherent source of motivation for seeking out and mastering optimal challenges” (ibid: 6). Relatedness is “the need to establish close emotional bonds and attachments with other people” (ibid). We discuss these in more detail and in relation to translation in the Section ‘Relevance to Translation Technology’ below.

Oulasvirta and Blom sound a warning about personalisation efforts in general, maintaining that too many attempts have been made at personalisation without regard to what people really want and “increasing availability of new features has coincided with underutilization of services, as well as degradations in usability and user acceptance” (ibid: 2).

In summary, then, it can be said that the motivation behind personalisation should be to increase the well-being of the system’s user by increasing autonomy, competence and relatedness.

User modelling and personalisation

Humans' conceptual understanding of the systems they interact with has advanced in recent decades, but there have been little advances in the development of user models by systems (Karat et al. 2004). Personalisation is not an easy task and techniques are still being sought that would allow for "…a future in which human-computer interaction is greatly enhanced through advances in the ability of technology to utilize personal information about users to realize better, more valuable interactions" (ibid. 2004: 9).

Generally speaking, personalisation is achieved through user modelling. User modelling can involve learning about 'interests' through online like/dislike votes, for example. Additionally, software can analyse web pages to determine what features caused the user to be (dis)interested, e.g. keyword extraction. Length of time spent on a page might also be used, as well as metadata in the HTML mark-up and number of click throughs from a page. In the Section'Relevance to Translation Technology', we propose how this modelling of 'interests' could be transferred to the translation context, for example by taking note of terminology look-up by the translator.

Munnelly et al. (2013) present an online user modelling model that has four phases: Guide, Explore, Reflect, Suggest. At first, the online system guides the user and learns from them as they explore the information in question. The user is then afforded a phase to reflect on and examine the user model produced by the system and to make changes (e.g. giving keywords more or less weight) to improve that model. The system subsequently uses the modelling information to suggest content that the user might be interested in exploring further. However, what actually matters for user modelling is not altogether clear and some tasks might require different inputs to build a user profile compared with others. This four-phase model is necessarily iterative as the user's interests and expertise evolve through exploration, both within the system and from external sources. According to Soltysiak and Crabtree (1998), user modelling for personalisation takes four aspects into consideration: Content, Acquisition, Privacy and Trust. We discuss these in relation to translation below.

A differentiation between user-adaptive (automatic) personalisation and user-driven (adaptable) systems is made (Oulasvirta and Blom 2008). The former is an automatic process, the latter is controlled by the user. Soltysiak and Crabtree (1998) suggested that user profiles should be acquired automatically with minimal user input. In contrast, Oulasvirta and Blom warn that “the need for autonomy and self-determination is seriously risked when personalisation, as a process that aims to enhance personal significance to the user, is driven by a computer instead of the user. In the worst case, the user is deprived of being the source of action, of being the author of decisions between action alternatives, and of the feeling of unpressured willingness to engage in an activity” (2008: 15). Therefore a balance needs to be maintained between implicit and explicit modeling such that accurate and timely models of the user can be assembled without overburdening the user.

Personalisation and Trust

Briggs et al. (2004) discuss personalisation and trust in the context of e-commerce and argue that the two are related and that trust is normally seen as a prerequisite for good personalisation. They suggest that the converse might also be true – good personalisation is a pre-requisite for trust building. Trust is an extremely difficult concept to work with since different factors are likely to be influential on trust at different times. The nature of trust depends on the ‘threat’. McKnight and Chervany (2001 – in Briggs et al. 2004) discuss the stages of trust building as (i) intention to trust and (ii) trusting activity. Briggs et al. (ibid.) add to this a third stage (iii) development of a trusting relationship, and highlight that very few studies focus on this type of trust development.

In the domain of e-commerce, scholars have generated a family of trust models; trust which supports online engagement “is influenced by perceived integrity and expertise, predictability or familiarity of content and reputation…A number of studies also highlighted the importance of interface factors (ease of use and functionality)” (Briggs et al. 2004: 43). In e-commerce studies, tailoring of information (selection of content according to your previous preferences, recognising you as a previous user etc.) was found to increase ‘credibility’, which is in itself seen as a factor of trust.

In the domain of Technology Enhanced Learning (TEL), research has been undertaken on Open User Models (Conejo et al. 2011, Kay 2008, Kobsa 2007), specifically in an attempt to promote scrutability (Bull and Kay 2007, 2013), a word coined to represent the user examining their user model. Open Learner Models usually present learner models through visualisations to support reflection and allows students to participate in the construction or modification of their personal model. It uses student responses to questions, number of attempts, and task response times, to build models of student competencies and levels of understanding. These models visualise learner competencies and levels of understanding, supporting reflection, assessment and monitoring (Bull and Kay 2013). The visualisations are inferred from learner interactions and have been shown to guide learning, help improve the performance of weaker students (Yousuf and Conlan 2015), and motivate stronger learners (Mitrovic and Martin 2007, Bull and Kay 2007). Open Learner Models and scrutability have typically given the user some control over the data modelled about them, rather than control over personalisation techniques directly.

Relevance to Translation Technology?

Here we consider how the field of personalisation and adaptation could be relevant to translation technology. We focus on the primary concepts alluded to above, i.e. the importance of context, user modelling, trust, motivation and well-being.

Context

García-Barrios et al. (2005) present a personalisation model that requires input from a modelling engine. The modelling engine, in turn, is divided into a User Model and other models that take the 'environment' into account. As discussed above, context in personalisation refers broadly to aspects of a situation that are internally represented in the computer. Context is a necessary and important aspect of all translation work too, but is known to be especially challenging for translation technology (Killman 2015). How then could personalisation engines take 'context' into account in translation? We propose that there could be a theoretical link between context modelling and translation specifications, as the latter are espoused by Hague et al. (2011). A specification defines what is expected in the translation task and includes parameters such as (but not limited to) audience, purpose and register. In a theoretical translation personalisation engine, models of the translation specifications could be built in order to control the output from the machine translation engine. Hague et al. (2011) discuss the importance of structured translation specifications, especially in the context of translator training and Melby more recently (personal communication: July

2016) proposes the association of structured translation specifications with points on a spectrum between fully automatic Machine Translation and full Human Translation. There are 21 standardised, hierarchical parameters proposed in the American Society for Testing and Materials' Standard Guide for Quality Assurance in Translation ASTM F2575-14 (American Society for Testing and Materials 2014), organised under the categories of Product, Process and Project. While it may not be feasible to operationalise all 21 specifications in a context model of a personalisation translation engine, it would be theoretically possible to adapt the engine according to some specifications, such as relative importance of productivity, accuracy, text type, domain and end user requirements. For example, translation of a text in the medical domain where accuracy is of very high importance for end users might have very high thresholds for quality estimation so that the translator is much less likely to erroneously accept an incorrect translation. In this context, a translator might even turn off the MT component and only work with exact or high fuzzy matches from a reliable TM. Conversely, a translation of a customer review for a new coffee maker which has to go on a website in 24 hours might have a lower threshold for quality estimation.

User Modelling

It was mentioned above that individuation of learning materials resulted in better outcomes in e-learning environments. The considerable array of translation process research that has been conducted in recent years suggests that, while translators often approach the translation task in similar ways, for example, by adopting an orientation, drafting and monitoring stage, their sub-processes and solutions are often quite individualistic. And yet difference in translation solutions does not mean that one solution is correct while the alternatives are incorrect. (This is perhaps why translation quality assessment is fraught with subjectivity and it proves difficult to even agree on a definition of translation quality (Melby et al. 2014, Fields et al. 2014, Koby et al. 2014). Just as personalisation in learning is put forward as being beneficial to learners, personalisation in computer-aided translation may also be beneficial. If we assume that translators might approach the task of translation in different ways, that they are likely to have different levels of tolerance for MT quality, and that this might depend on the task at hand (i.e. context), including the language pair they are working with, then we can also assume that personalisation according to these requirements, likes and tolerances would be a good idea, in principle at least.

User modelling would be required so as to personalise technology according to individual translator profiles. As mentioned previously, user modelling for personalisation takes four aspects into consideration: Content, Acquisition, Privacy and Trust (Soltysiak and Crabtree 1998). We consider here how these might be applied to modelling for a translator.

In the realm of personalisation and user modelling, content is seen to reflect users' interests (e.g. what they search for, how often, which pages they browse or ignore etc.) and acquisition involves learning from this type of information. For translator modelling, a personalisation engine could draw on the information sources searched (e.g. encyclopedic, specialised content from expert content producers, monolingual, bilingual or multilingual websites and dictionaries etc.), how frequently such resources are checked, for how long, how much double or triple checking is done for terminology, collocations etc. In this way a personalised translation engine could learn about the translators' 'interests', which are, of course, driven by their clients, and also about the keywords (terms) they need to understand and which mono- and multilingual resources they rely on and trust.

Acquisition of information for personalisation could additionally be acquired by learning not only from the translators' search activities but also from the number and types of edits for TM matches and MT suggestions. Given the considerable research to date involving keyboard logging for translation process research, the implementation of these features in commercial TM tools (e.g. iOmegaT, MateCAT), as well as automatic post-editing and machine learning research in Natural Language Processing in general, it is not unrealistic to suggest that these techniques could be deployed to acquire information for translator modelling. Furthermore, eye tracking has been used as a method for measuring cognitive load in language processing in general, and in translation and post-editing research specifically. In principle, eye tracking could also be used to determine what information a translator pays most attention to and what aspects of a text demand most processing effort. Such information could then be used to tune a personalised translation engine. We can reasonably expect eye tracking technology to eventually be embedded in our computing hardware, making this even more of a reality in the (near) future.

User modelling using such techniques inevitably raises concerns about ethics and privacy. Individual translators would have to give consent to be profiled in this way. The translator would also have to be convinced that divulging task activity would have a high probability of leading to a better technological experience in the longer term and that it would not interfere in a detrimental way in their normal process, which is, after all, one on which they rely to make a living. Being monitored in such a way potentially exposes a lack of knowledge or weakness in the individual translator, as well as providing evidence of good practice and expertise. The privacy of information acquired for user modelling would therefore need to be guaranteed.

The fourth aspect listed by Soltysiak and Crabtree is trust, which, alongside privacy, is a considerable issue and is dealt with in more detail below.

In the previous section, we mentioned Oulasvirta and Blom's (2008) differentiation between user-adaptive (automatic) personalisation and user-driven (adaptable) systems. Automatic modelling is still some ways off, it would seem, and this may especially be the case for multi-faceted expert tasks such as translation. Supervised learning of individual profiles might be more realistic in the shorter term. In this scenario, the user would 'teach' the system about preferences and needs. The questions that emerge here for translation is whether data like edit distance, fixations, number of searches etc. could be used as 'semi-automated' supervised learning of translator profiles and whether these could be linked with the contextual models (structured translation specifications) mentioned earlier. Would this data be too 'noisy' to be meaningful in the generation of a personalised computer-aided translation tool? Should the personalisation model also factor in the translation revision cycle, both self-revision and third party revision (e.g. quality scores, reviewer feedback, number of comments, disagreements etc.)?

A final question that arises is whether there should be one profile per translator or many? A translator might have a profile when working on one topic for one client and a different profile for another, depending on task specifications. The profile(s) might also change over time. Automatic detection of changes in profile over time may be useful.

Trust

The concept of trust has recently been linked with translation in general and translation technology in particular, especially MT (Pym 2013, 2015, Cadwell et al. 2017, Abdallah and Koskinen 2007).

Translators will make use of metadata in TM systems to decide on their trust levels for a TM match. A translator known to them will, generally, be awarded more trust than one who is not known. The fuzzy match score and markup of differences will also contribute to levels of trust. In the context of MT, translators have been known to report that they do not trust the output and will check it more thoroughly if they know it has been generated by an MT system (Bundgaard 2017). Quality estimates from the MT system should, in theory, guide the trust levels but this technology is too new at this point and translators are likely to be quite sceptical of a machine's rating of its own output (Moorkens and O'Brien 2016).

A personalisation system could learn, through user modelling, about the trust levels and techniques used to establish trust by an individual translator. Moreover, personalisation could offer a scrutable trace to the translator to support their understanding of why a particular translation is being suggested. For example, length of time spent on an MT suggestion compared with a TM suggestion (if the two are clearly differentiated) could be taken as an indicator of trust, as could edit distance data, i.e. how much the suggestion has been edited. Information about the origin of the MT suggestion (whether from System A or System B) could be factored in. Trust levels could also be estimated using gaze data from an eye tracker, e.g. how often is the source segment re-read, what metadata does the translator look at, how often does she check the glossary or external lexicographical resources, does she take the confidence score into account when making editing decisions, or does she largely ignore that metadata?

Apart from learning about what drives trust in a translator, reliable personalisation techniques could, in principle, also lead to higher trust among translators. If the personalisation techniques are seen to be successful, the translator will trust the personalisation engine more. Until that success is experience, trust levels for the personalisation translation engine are likely to be low.

Motivation and Well-Being

Earlier we alluded to the fact that the introduction of TM technology was top-down and mostly treated as a technical change process and not a socio-technical change process (Olohan 2011). With increased technologisation, the risks of negative impacts on translator motivation and well-being are high. As already discussed, there is evidence that while translators mostly find TM tools very helpful in their jobs, they are still, twenty or more years on, dissatisfied with certain aspects of them. They fear MT, its impact on their language and translation skills and their creativity, and also fear being ‘reduced’ to merely a post-editor.

The question that emerges then is whether personalisation and adaptation of translation technology has the potential to contribute positively to translator motivation and well-being? Could CAT be developed further as a “personal agent”, i.e. “software capable of operating autonomously in order to provide timely and relevant information for an individual” (Soltysiak & Crabtree, 1998: 110)?

Through personalisation and adaptation techniques, could CAT positively influence the three dimensions listed above: Autonomy, Competence and Relatedness? If the translator is pressured to use the technology then feelings of autonomy are likely to be compromised. In the emerging scenario where translation technology is becoming more pervasive, an important question is how translator autonomy can be preserved while also interacting with CAT tools. Having control over how the technology serves the individual translator’s working methods through a personalisation engine might contribute to feelings of autonomy.

If using CAT tools gives translators a sense of loss of competence, this will surely affect motivation and well-being. A challenge for the personalisation of translation technology, then, is to limit this effect and to create a personalised system that increases the sense of competence rather than having the opposite effect. For example, if the system learns about the translator’s individual tolerance for MT errors, it can customise the situations and frequency with which MT-generated text is presented to the translator. In this way, the translator may have less of a sense of being ‘reduced’ to ‘fixing’ errors produced by technology and the sense of being able to produce a fit-for-purpose text in a timely manner may be increased. Moreover, scrutability and control of the user model may foster more nuanced control over how CAT works for individual translators.

In one respect, CAT can be seen as increasing relatedness as it is a form of collaborative translation, where the translator reuses another translator’s suggestion, or benefits in real time from a shared online TM. Adaptive and interactive MT might also offer a sense of ‘relatedness’. However, interaction with a machine, especially in circumstances where the machine is learning and benefitting from the human activity could also have the opposite effect. Another challenge for personalisation is how to increase relatedness between the translator and his tools? A focus on creating software that serves the individual and therefore increases relatedness is desirable.

Conclusion and a Research Agenda

This chapter reflected on some of the major boundary shifts that have occurred to date in the domain of translation technology and on the impact that has had on the translator, the process and the product. It then introduced the concepts of personalisation and adaptation and described how they have been deployed in e-commerce and e-learning, and the most important inputs: context, motivation, user modelling, trust and well-being. We turned our attention then to considering how personalisation might be relevant to translation technology and the translator, suggesting that a theoretical personalised translation engine could take account of concepts in translation studies, e.g. translation specifications, measurement techniques, e.g. keyboard logging, eye tracking, edit distance, and research innovations, e.g. quality estimation, to build a personalised translation technology engine that would serve to maintain autonomy, competence and relatedness, as well as motivation and well-being for professional translators in an increasingly technologised profession.

This chapter represents initial ideas and explorations into how personalisation might be used to the advantage of translators. As mentioned previously, personalisation is not an easy technique and so considerable research would be required to determine if and how it might be useful in the field of translation. We conclude here by highlighting some research directions that need to be addressed. What type of translation data might be of most use in a personalised translation engine? Edit distance, fixation, search or temporal data, or some combination of these? And would the usefulness of the data depend on the stage of translation (e.g. drafting versus revision)? How might translation specifications be formalised in a personalised translation engine? How willing are translators to engage in a user-driven personalisation process? Should there be more than one profile per translator? Would personalisation contribute to feelings of autonomy, competence and relatedness? How could personalisation build trust among translators, of both the personalisation process and of the translating machine?

These questions suggest a challenging research agenda for the future. The alternative option of continuing with a generic technology that tries to fit all contexts and suit all translators is not very compelling.

References

Abdallah, Kriistina, & Koskinen, Kaisa 2007. Managing trust: Translating and the network economy. *Meta* 52 (4), 673-687.

ASTM F2575-14 2014. *Standard guide for quality assurance in translation*. ASTM International, West Conshohocken, PA. www.astm.org. (accessed 01 July 2017).

Billsus, D., Pazzani, M. J., & Chen, J. 2000. A learning agent for wireless news access. In D. Riecken, D. Benyon & H. Lieberman (eds) *Proceedings of the 5th international conference on intelligent user interfaces*. ACM, 33-36.

Bull, S., & Kay, J. 2007. Student models that invite the learner in: The SMILI:() Open learner modelling framework. *International Journal of Artificial Intelligence in Education* 17 (2), 89-120.

Bull, S., & Kay, J. 2013. Open learner models as drivers for metacognitive processes. In R. Azevedo & V. Aleven (eds) *International handbook of metacognition and learning technologies*. New York: Springer, 349-365.

Bundgaard, Kristine 2017. *(Post-)Editing: A workplace study of translator-computer interaction at Textminded Danmark A/S.* PhD Dissertation, Aarhus University.

Briggs, Pamela, Simpson, Brad & de Angeli, Antonella 2004. Personalisation and trust: A reciprocal relationship? In C.M. Karat, J. Blom & J. Karat (eds) *Designing personalized user experiences in eCommerce*. Human-Computer Interaction Series 5. New York: Kluwer, 39-56.

Brusilovsky, P. 2008. Adaptive navigation support for open corpus hypermedia systems. In W. Nejdl, J. Kay, P. Pu, & E. Herder (eds) *Adaptive hypermedia and adaptive web-based systems*. AH 2008. Lecture Notes in Computer Science, vol 5149. Berlin/Heidelberg: Springer, 6-8.

Cadwell, Patrick, Teixeira, Carlos, O'Brien, Sharon 2017. Resistance and accommodation: Factors for the (non-)adoption of machine translation among professional translators. *Perspectives: Studies in Translation Theory and Practice* 26(3), 301-321.

Conejo, R., Trella, M., Cruces, I., & Garcia, R. 2011. INGRID: A web service tool for hierarchical open learner model visualization. In L. Ardissono & T. Kuflik (eds) *Advances in user modeling*. UMAP 2011. Lecture Notes in Computer Science, vol 7138. Berlin/Heidelberg: Springer, 406-409.

Conlan, O., Wade, V., Bruen, C., & Gargan, M. 2002. Multi-model, metadata driven approach to adaptive hypermedia services for personalized e-learning. In P. de Bra, P. Brusilovsky, & R. Conejo (eds) *Adaptive hypermedia and adaptive web-based systems*. New York: Springer, 100–111.

Dara, Aswarth, Dandapat, Sandipan, Groves, Declan, & van Genabith, Josef 2013. TMTprime: A recommender system for MT and TM integration. In *Proceedings of NAACL HLT 2013, Demonstration Session*, 10-13.

De Bra, P., Houben, G. J., & Wu, H. 1999. AHAM: a Dexter-based reference model for adaptive hypermedia. In J. Westbomke, U. Wiil, J. Leggett, K. Tochtermann, J. Haake (eds) *Proceedings of the tenth ACM Conference on hypertext and hypermedia: returning to our diverse roots*. ACM, 147-156.

Dragsted, Barbara 2005. Differences across levels of expertise and difficulty. *Target* 17 (1), 49-70.

Fields, Paul, Hague, Daryl, Koby, Geoffrey, Lommel, Arle, Melby, Alan 2014. What is quality: A management discipline and the translation industry get acquainted. *Revista Tradumàtica* 12, 404-412.

Forcada, Mikel 2017. Making sense of neural machine translation. *Translation Spaces* 6 (2), 291-309.

Foster, George, Isabelle, Pierre, & Plamondon, P 1997. Target-text mediated interactive machine translation. *Machine Translation* 12 (1–2), 175–194.

Foster, George, Langlais, Pierre, & Lapalme, G. 2002. User-friendly text prediction for translators. In *Proceedings of the 2002 conference on empirical methods in natural language processing 10*, 148–155.

Green, Hannah, Facer, Keri, Rudd, Tim, Dillon, Patrick, & Humphreys, Peter 2005. Futurelab: Personalisation and digital technologies. Research report. HAL Id: hal-00190337. https://telearn.archives-ouvertes.fr/hal-00190337. (accessed 7 July 2017).

Garlatti, S., Iksal, S., & Kervella, P. 1999. Adaptive on-line information system by means of a task model and spatial views. In P. Brusilovsky & P. de Bra (eds) *Proceedings of the second workshop on adaptive systems and user modeling on the* *world wide web*, 59-66.

García-Barrios, Victor Manuel, Mödritscher, Felix & Gütl, Christian 2005. Personalization versus adaptation? A user-centered model approach and its application. In *Proceedings of I-KNOW*, 120-127.

Göker, Ayse & I Myrhaug, Hans 2002. User context and personalisation. In Mehmet H. Göker & Barry Smyth (eds) *Workshop on case based reasoning and personalization*, 6th European Conference on Case Based Reasoning ECCBR. Aberdeen, Scotland (no page numbers).

Hague, Daryl, Melby, Alan, Zheng, Wang 2011. Surveying translation quality assessment: A specification approach. *The Interpreter and Translator Trainer* 5 (2), 243-267.

Hampson, C., Lawless, S., Bailey, E., Hillemann, E., Steiner, C., & Conlan, O. 2014. Metadata-enhanced exploration of digital cultural collections. *Journal of Metadata, Semantics and Ontologies* 9 (2), 155-167.

Höök, K. 1998. Evaluating the utility and usability of an adaptive hypermedia system. *Knowledge-Based Systems* 10 (5), 311-319.

Joerding, T. 1999. A temporary user modeling approach for adaptive shopping on the Web. In P. Brusilovsky & P. de Bra (eds) *Proceedings of second workshop on adaptive systems and user modeling on the world wide web*, 75-79.

Karat, John, Karat, Clare-Marie, & Brodie, Carolyn 2004. Personalizing interaction: Directions for HCI Research. In C.M. Karat, J. Blom & J. Karat (eds) *Designing personalized user experiences in eCommerce*. Human-Computer Interaction Series 5. New York: Kluwer, 7-17.

Kay, J. 2008. Lifelong learner modeling for lifelong personalized pervasive learning. *IEEE Transactions on Learning Technologies* 1 (4), 215-228.

Killman, Jeffrey 2015. Context as Achilles' heel of translation technologies: Major implications for end-users. *Translation and Interpreting Studies* 10 (2), 203-222.

Koby, Geoffrey, Fields, Paul, Hague, Daryl, Lommel, Arle, & Melby, Alan 2014. Defining translation quality. *Revista Tradumàtica* 12, 413-420.

Knutov, E., De Bra, P., & Pechenizkiy, M. 2009. AH 12 years later: a comprehensive survey of adaptive hypermedia methods and techniques. *New Review of Hypermedia and Multimedia* 15 (1), 5-38.

Kobsa, A. 2007. Generic user modeling systems. In P. Brusilovsky, A. Kobsa & W. Nejdl (eds) *The adaptive web*. Berlin/Heidelberg: Springer, 136-154.

Koehn, Philip 2010. *Statistical machine translation*. Cambridge University Press, Cambridge: UK.

Kurzweil, Ray 2010. The singularity is near. London: Duckworth Overlook.

Lawless, Séamus, Conlan, Owen, & Hampson, Cormac 2016. Tailoring access to content. In Susan Schreibman, Ray Siemens and John Unsworth (eds) *A new comparison to digital humanities*. New Jersey: John Wiley, 171-184.

LeBlanc, Matthieu 2013. Translators on translation memory (TM): Results of an ethnographic study in three translation services and agencies. *Translation and Interpreting* 5 (2), 1–13.

Melby, Alan 2016. Professional translation in a pre-singularity world. Oral presentation at the *Asling Translation and Computer 38 Conference*, November 2016. London: UK.

Melby, Alan, Fields, Paul, Hague, Daryl, Koby, Geoffrey, & Lommel, Arle 2014. Defining the landscape of translation. *Revista Tradumàtica* 12, 392-403.

Mellinger, Christopher & Shreve, Gregory M. 2016. Match evaluation and over-editing in a translation memory environment. In R. Muñoz Martín (ed) *Reembedding Translation Process Research*. Amsterdam: John Benjamins, 131-148.

Mitrovic, A., & Martin, B. 2007. Evaluating the effect of open student models on self-assessment. *International Journal of Artificial Intelligence in Education* 17 (2), 121-144.

Moorkens, Joss & O'Brien, Sharon 2016. Assessing user interface needs of post-editors of machine translation. In Dorothy Kenny (ed.) *Human issues in translation technology - IATIS yearbook 2016*. London, UK: Routledge, 109-130.

Munnelly, Gary, Hampson, Cormac, Conlan, Owen, Bailey, Eoin & Lawless, Seamus 2010. Personalising the cultural heritage experience with CULTURA. In *Proceedings of ENRICH, SIGIR Workshop*, August 1-2, 2010, Dublin, 39-42.

O'Brien, Sharon, Ehrensberger-Dow, Maureen, Hasler, Marcel & Connolly, Megan 2017. Irritating CAT tool features that matter to translators. In T. Paulsen Christensen, M. Flanagan & A. Schjoldager (eds) *Special Issue of Hermes: Journal of Language and Communication in Business on Translation Technology Research* 56, 145-162.

O'Hagan, Minako 2013. The impact of new technologies on translation studies: A technological turn? In Carmen Millán and Francesca Bartrina (eds) *Routledge Handbook of Translation Studies*, 503-518.

Olohan, Maeve 2011. Translators and translation technology: The dance of agency. *Translation Studies* 4 (3), 342-357.

Oulasvirta, A. & J. Blom 2008. Motivations in personalisation behaviour. *Interacting with computers* 20 (1), 1-16.

Pym, Anthony 2015. Translating as risk management. *Journal of Pragmatics* 85, 67-80.

Pym, Anthony 2013. Translation skill-sets in a machine translation age. *Meta* 58 (3), 487-503.

Pym, Anthony 2011. What technology does to translation. *The International Journal for Translation & Interpreting Research* 3 (1), 1-9.

SDL 2016. *Research study 2016: translation technology insights*. http://www.sdl.com/solution/language/translation-productivity/research-2016/ (accessed 04 January 2017).

Soltysiak, S.J. & Crabtree, I.B. 1998. Automatic learning of user profiles – towards the personalisation of agent services. *BT Technology Journal* 16 (3), 110-117.

TAUS 2014. *Machine translation market report*. https://www.taus.net/think-tank/reports/translate-reports/mt-market-report-2014 (accessed 01 June 2017).

Vasconcellos, Muriel 1985. Management of the machine translation environment: Interaction of functions at the Pan American Health Organization. In Veronica Lawson (ed) *Tools for the trade, Proceedings of Translating and the Computer 5*, London: Aslib, 115-129.

Vasconcellos, Muriel & León, Marjorie 1985. SPANAM and ENGSPAN: Machine translation at the Pan American Health Organization. *Computational Linguistics* 11 (2-3), 122-136.

Vashee, Kirti 2016. *The emerging world of neural net based MT*. http://kv-emptypages.blogspot.ie/2016/06/the-emerging-world-of-neural-net-driven.html. (accessed 01 June 2017).

Yousuf, B., & Conlan, O. 2015. VisEN: Motivating learner engagement through explorable visual narratives. In G. Conole, T. Klobučar, C. Rensing, Konert, J., & É. Lavouré (eds) *Design for teaching and learning in a networked world.* Proceedings of EC-TEL 2015. Berlin/Heidelber, Springer: 367-380.

Zetsche, Jost 2016. Lilt: Translation environment tool of a different kind. *Multilingual* 157, January/February 2016, 15-17.