

THE ETHICS OF MACHINE TRANSLATION

Dorothy Kenny

Introduction

This paper grew from a presentation at the annual conference of the *New Zealand Society for Translators and Interpreters* in Christchurch in July 2010. The theme of the conference was ‘service, value and ethics’, which gave me an opportunity to reflect on a number of issues that had been troubling me for some time. Those issues revolved (and still revolve) around the development and use of machine translation systems, and their implications for the teaching of translation. The malaise that I was experiencing at the time could not be attributed to deep-seated antagonism against machine translation. I had, after all, completed a postgraduate degree in machine translation (admittedly in the early 1990s) and had been happily teaching budding translators about the technology for many years. I enjoyed participating in a community in which machine translation researchers and translation scholars could understand each other and work together,¹ and I appreciated that machine translation could act as a test bed for lexical and grammatical formalisms that attempted to capture the complexity of natural languages. And even if such formalisms were not always successful (because it is nearly impossible to capture in this way in all the knowledge one needs to translate real texts, unless one is dealing in a very constrained specialized domain), the approach had many notable successes. In short, I could live with a technology that I understood, whose developers I could talk to, and whose limitations were well known. I could live with a technology that co-existed with human translation without really rivalling human translation. More importantly, I could live with a technology that was well differentiated from human translation: machine translation was not human translation, and human translation was not machine translation. To put it even more briefly, I could live with rule-based machine translation. But everything changed with the advent of statistical machine translation.

In this paper I first describe these two branches in machine translation research. I then go to discuss why the second of these, statistical machine translation, can cause such malaise among translation scholars. As some of the issues that arise are ethical in nature, I stop to ponder what an ethics of machine translation might involve, before considering the ethical stance adopted by some of the main protagonists in the development and popularisation of statistical machine translation, and in the teaching of translation.

Machine Translation Paradigms

As suggested above, rule-based machine translation (RBMT) uses formal grammars and lexicons to represent the grammatical structures and words that are ‘allowed’ in a given source and target language. In its most typical instantiation it also uses bilingual ‘transfer’ rules to map between the grammatical structures and words of the two languages. Rule-based systems are generally considered expensive to develop as they require trained computational linguists to write the grammatical and lexical rules, and they often encounter knowledge bottlenecks beyond which it is difficult to improve the performance of the system. These factors are sometimes considered to have stalled

¹ See Kenny and Way 2001 (1) for an example of a paper written with both trainee translators and trainee computational linguists in mind.

progress in RBMT, although RBMT systems continue to be developed to this day. RBMT is exemplified by the open-source system Apertium, and earlier iterations of Systran.²

Statistical machine translation (SMT), exemplified by a number of well-known systems, including Google Translate, Microsoft Translator, Language Weaver and Asia Online, on the other hand, does without such expensive grammars and lexicons. Rather, it seeks to ‘learn’ probable translations of source language strings of varying lengths (for example, single words, or strings of two or three words, etc.) from an already existing parallel corpus of source texts and their human translations. All other things being equal, the bigger the parallel corpus an SMT system is trained on, the better the statistical model it can build of translation between the two languages in question. SMT thus relies on human translation for its training data, but it also relies on human translation for its legitimacy: the reason developers of SMT systems use parallel corpora to train their systems is because such corpora are assumed to contain good answers to translation problems; and they are assumed to contain good answers precisely because they contain translations performed by human beings.

There is more to SMT than the statistical translation models learned from existing translations, of course. SMT systems also use monolingual models of the target language to work out whether a candidate translation is actually a good string in the target language, and there is considerable ingenuity involved in generating candidate translations (using translation and target-language models, as well as various other ‘features’) and subsequently selecting the most probable one. A full discussion of how SMT works is beyond the scope of this paper; the reader is referred instead to Hearne and Way 2011 (2), who provide a much more sophisticated, yet still accessible introduction to the area. What is important for current purposes is that SMT relies on the re-use of human translations. It is thus a technology in which human and automatic translation intersect, which itself can be a discomfiting thought for those who were previously comforted by clear lines of demarcation between the two areas. It is also a technology that relies on the ingenuity of both human translators (who produce vital data) and statistically-minded computer scientists (who work out clever ways of using these and other data), and both sets of protagonists might expect to be acknowledged in discussions of SMT.

The rise of SMT

Although machine translation has been around since the 1950s, it has received a new lease of life in the last two decades, not least because of the advent of SMT. Since the proof of concept stage, which started with experiments at IBM in the late 1980s, SMT has grown to dominate machine translation research, and the consensus is that performance continues to improve. SMT is considered especially successful between languages that do not exhibit vastly different structures from each other and for which adequate training data is available. Indeed, SMT has benefited from the increased availability of aligned parallel corpora, which have been built, for example, from data created by bilingual government sources, multilingual organizations like the United Nations and the European Union, and multinational enterprises who have built up large translation memories since the mid 1990s. SMT has also benefited from

² Systran is now better described as a ‘hybrid’ system, as its later instantiations incorporate statistical processes.

increasing computer power, which is vital in the training of large-scale SMT systems, as well as the growing popularity of shared resources, including corpora and open-source code. At the same time the number of ways in which machine translation can be delivered to consumers (often at no financial cost) has increased: a computer user can request a machine translation by pressing a button in a Microsoft Office application, or through his or her search engine, or through a widget embedded in a web page, for example. And professional translators have the additional option of linking to a machine translation system directly from their translation memory tool, an option considered useful if the translation memory currently in use does not return a high-value fuzzy match.

All of these factors (along with the much vaunted ‘explosion’ of digital content) have combined to produce a kind of perfect storm for machine translation, so much so that by 2006, Jaap van der Meer, one of the leading commentators in the area was claiming that despite a number of false dawns over the decades, ‘Everything is playing along to make machine translation technology successful this time’ (3). What this means is that machine translation is more than ever a force to be reckoned with, for producers, consumers and translators of digital content alike.

Ethics and Translation

Little has been written to date on the ethics of machine translation, although ethical issues have been approached from a variety of angles in the wider area of translation studies. I refer here to just two translation theorists whose work both gives an overview of contemporary thinking and moves the discussion on somewhat: Chesterman 2001 (4) and Baker 2011 (5). The discussion then moves to contributions more directly concerned with machine translation.

Chesterman (*ibid.*) recognises four models underlying scholarship in the area of translation and ethics: the first focuses on representation (a concern here might be how the target text represents the source text or author); the second on service (here we might ask how translators can act ethically vis-à-vis their clients); the third on communication (how does translation operate in encounters with the Other?); the fourth on norms (do translators act as they are expected to?). Chesterman goes on to argue that the different models of ethics that underly these approaches are often incompatible with each other, and that each is necessarily partial, covering different ground to the others. He also points to some gaps in the approaches he surveys, one of which relates to the world’s responsibility towards translators. This, he argues, might also be considered ‘to belong to a general ethics of translation and translatorial behaviour’ (*ibid.*:143). Chesterman ultimately draws on virtue ethics to explore an alternative way of looking at ethics and translation. Following MacIntyre 1981 (6), he defines a virtue as ‘an acquired human quality that helps a person strive for excellence in a practice’ (Chesterman *ibid.*:145), and goes on to suggest that in order to make the best ethical decisions, the most important virtue that a translator can possess is the desire to make the right decision: ‘the translator must *want* to be a *good* translator, must strive for excellence in the practice of translation’ (*ibid.* emphasis in the original). Chesterman’s contribution is particularly useful here because it points up the partial nature of many approaches to translation and ethics and some of the gaps that arise in these approaches. It also reminds us of two other important points: translators themselves can be affected by the ethical (or unethical) behaviour of others

(although Chesterman does not pursue this idea), and translation ethics can be defined in terms of translators themselves striving for excellence in the practice of translation.

In her more recent contribution, Baker uses some of the broad strands of moral philosophy (utilitarianism, Kantian rule-based ethics, etc.) to frame a discussion of the ethical decisions that translators and interpreters face in their work. Baker refers to the distinction that some commentators make between ‘ethics’, or codes of conduct that constrain the behaviour of members of a given profession (such codes roughly align with Chesterman’s ethics of service), and ‘morality’, which is often construed as individual and pertaining generally to decisions made in the daily flux of life. Baker ultimately rejects this distinction, however. She also rejects the idea that ethics apply only in instances where we are faced with extraordinary moral choices. Following Cheney *et al.* (7), Baker argues that ‘ethics is about the stream of life rather than just its turbulent moments’ (Cheney *et al.* 2010:237, in Baker 2011: 277), and that ‘many default choices that do not necessarily give rise to conscious decision-making can have important ethical implications’ (Baker *ibid.*). It is this reminder of the ethical import of decisions we take in our every-day lives (or the routine decisions we take in our working lives) that is very valuable for current purposes.

Both Chesterman and Baker are primarily concerned with what translators (and interpreters) do however, rather than how translators are affected by ethical decisions taken by other parties. When it comes to machine translation, as we will see below, we need to factor in other decision-making agents.

Ethics and Machine Translation

As already indicated, little has been published on ethics and machine translation. A notable exception is Melby and Warner’s 1995 (8) use of the philosophy of Emmanuel Levinas to question the very possibility of machine translation in all but the above-mentioned constrained specialized domains (or what they call ‘utterly boring worlds’). For Melby and Warner, dynamic general language is based at least partly on an ethical stance that acknowledges the ‘Otherness’ of interlocutors who are outside our control and whose inner life and individuality we recognise. In order to communicate with others, we must have agency, which involves the capacity to make real choices for which we take responsibility, and we must also regard our interlocutors as having agency. In a passage that summarizes the argument about agency, Melby 1995 (9) writes:

Unless we regard others as agents, just like us, who in turn regard us as agents, then many key notions that are a basis for general vocabulary become meaningless. Without this interacting agency, there is no responsibility, no empathy or indifference, no blame, and no gratitude. So much becomes missing from language that what is left can be described as a technical domain and handled by a computer...Without agency, we are reduced to the status of machines and there is no dynamic general language. Without dynamic general language, we would translate like computers and there would be no truly human translation as we now know it. Thus lack of agency is one factor that keeps computers from translating like people.

Note that Melby does not argue that machine translation is impossible. It is possible for computers to handle language in ‘technical domains’. Lack of agency, however,

sets the boundaries to the possibility of computers translating like humans in dynamic, general domains. To use one of Melby's (*ibid.*) examples, when translating a French menu, a human translator might stop to think that an English speaker in France would appreciate being told that a *steak tartare* is served entirely raw, even if this information is not contained in the original text (because French people might be assumed to know this already). Such a translator would be aware of differences in material culture, and would be able to empathize with the English speaker who might choose to avoid the dish, given more information. Because computers do not have agency, they cannot make such decisions, and so their translations are necessarily constrained compared to those of humans.

As should be clear from the above, Melby and Warner (1995) are concerned with the ethics of *communication*, rather than any of the other models of ethics identified by Chesterman and referred to above, and they are particularly concerned with the role of agency in communication. It is worth noting, however, that Melby and Warner's book predates the current ascendancy of SMT.³ Given that SMT relies to large extent on existing human translations, it is quite possible for an SMT system to generate a translation that reflects the kind of sensitivity to the Other evidenced in good human translation. Indeed, Franz Och (11), one of the brains behind Google Translate, gives such an example in a Google 'TechTalk'. Och explains how, when evaluating output of Google Translate, engineers at Google discovered that the system had translated 'Heath Ledger' into Spanish as 'Tom Cruise'. The explanation for the slip was simple: in the parallel corpus used to train the system, there was indeed an instance where a human translator had translated 'Heath Ledger' as 'Tom Cruise'. The translator had done this because Heath Ledger was mentioned in the original text merely as an example of a famous male movie star. In the case in question, the translator reasoned that the target language audience (in Argentina) was unlikely to recognise the name and so substituted that of a different male movie star, whose name they would recognise. Och calls such examples 'very interesting and fancy reformulations' and describes the Tom Cruise solution as 'a very good translation'. Whether the translation should have been reused in a different context is another story, but examples like this show how we can find clear traces of human translator agency (and in particular a human's ability to empathize with imagined readers) in contemporary SMT.⁴ This does not mean that the SMT system has agency, however. SMT systems make decisions, but they do so using statistical models and algorithms designed to select, as efficiently as possible, the 'best' option from a large number of hypotheses built using those models. They do not exercise will to make ethical choices that take into account the otherness of target language readers. Having said that, the intersection of human agency with SMT remains an interesting, though largely unexplored field.

A recent paper by Drugan and Babych 2010 (12) also merits mention here. Drugan and Babych are concerned with the ethical issues that arise when translation resources are shared. Such resources include parallel corpora created from source texts and their translations published by governments and international organizations, and used to train SMT systems. In such cases, according to Drugan and Babych, use of parallel

³ Although Marcu and Melby 2006 (10) ask some thought-provoking questions about data-driven MT (of which SMT is by far the most widespread type).

⁴ In fact, all SMT based on human translations by definition contains traces of human agency; these traces are just not usually as clear as in the Heath Ledger case.

resources ‘does not involve serious ethical issues’ as the data in question are not confidential, they are produced using public funds, ‘and from the outset, it is intended they will be available for everyone to consult and use’ (2010:4). But in other cases where translation memories are shared, or services such as Google Translator Toolkit retain source texts and their translations to assist in the training of the Google Translate SMT system, then significant issues of confidentiality arise (as the texts in question might contain sensitive personal or commercial information, for example). Not only that, but people who reuse translations through Google Translate cannot be sure that those translations have been shared by their rightful owners in the first place, and nor can they acknowledge the anonymous translators whose work they are reusing. At the heart of Drugan and Babych’s contribution is an understanding of ethics that is consistent with the ‘professional ethics’ or ‘ethics of service’ models identified above. Perhaps unsurprisingly then, they argue that the professionals codes of ethics used by translators’ associations can provide guidance on how to deal with some of the issues that arise in the sharing of translation resources.

The above contributions represent a good start in the formation of an ethics of machine translation. They address fundamental issues of communication between agents and begin to broach the idea that powerful interests who share translation resources might have some ethical obligations towards those whose intellectual production is re-consumed in the process. But like the models surveyed by Chesterman, they are necessarily partial, and do not yet take into account other stakeholders in machine translation scenarios. In the next section I attempt to broaden our view of the ethics of machine translation, and to do so in a way that addresses the particular issues that cause the above-mentioned malaise for a translation teacher.

Towards a Broader Ethics of Machine Translation

Given the vast number of sometimes incompatible ways in which we can think about ethical actions (or living morally), most moral philosophers agree to a minimum conception of morality, on which basis further discussion can take place. This minimum conception is summarized by Rachels and Rachels (13) as follows:

Morality is, at the very least, the effort to guide one’s conduct by reason—that is, to do what there are the best reasons for doing—while giving equal weight to the interests of each individual affected by one’s decision.
(2010:13)

Seen in this light, ethical decisions are reasoned decisions that take others into account. If we ask who those ‘others’ might be in situations in which SMT is used (who is ‘other’ of course depends on our own position, but we will try to cover multiple possibilities here), then we might conclude that an ethics of MT involves, among others:

- developers of SMT systems
- commissioners and consumers of machine translations (in cases where machine translation is used for assimilation purposes only, the consumer and the commissioner coincide)
- source text writers
- translators, whose production is used to train SMT systems, and who at the same time risk being displaced by SMT

- post-editors, who bring machine translation output up to a standard that is fit for purpose
- trainee translators, who are entering a profession that is in flux
- translation teachers, who need to respond to a changing environment and help prepare students for a sustainable and rewarding working life

It would be impossible to discuss in full here all the ethical considerations that arise if we take each of these stakeholders into account; what follows is thus necessarily partial (like other approaches addressed above). Given the exploratory nature of this paper, the discussion is also likely to scratch only the surface of the issues involved, but a discussion of the ethics of machine translation is urgently needed, and one has to start somewhere.

We might start our discussion by asking whether SMT developers and other protagonists take into account the interests of translators. This is a broad question. It might be recast as a series of narrower questions related to the visibility of translators or the sustainability of the translation profession. At a very basic level we might ask whether developers acknowledge translators as a source of data?

Some prominent pronouncements from commercial sources do indeed indicate the human provenance of the data on which SMT systems are trained. The developers of Google Translate, for example, indicate in their online support documentation that:

we feed the computer...aligned text consisting of examples of human translations between the languages. (<http://translate.google.com/support/>)

But while some Microsoft sources indicate that there is human input in SMT, other ones (like Doug Thomas' blog at <http://blogs.msdn.com/b/translation/>) do not.

Within the academic community of SMT developers it is perhaps fair to say that the role of translators in creating vital data has been mostly downplayed or ignored, although Way and Hearne 2011 (14) are adamant that 'the role of the translator in SMT is a crucial one: they provide *all* the knowledge upon which our models are based'. (In this and other commentaries on SMT, Andy Way has been consistent in his recognition of the translator's input.)

There are more subtle ways of making translators visible and invisible, however. One is based on the metaphors that are used to describe the process of SMT, and is already operational in the description of SMT provided by Google and reproduced above. Here translation is conceptualised as food ('we feed the computer...human translations...'). Elsewhere translations are conceptualised, like other types of 'big data', as a crop to be harvested (see, for example, Quah's (2006:79) contention that the goal in SMT is 'to harvest a list of possible translation equivalents for a new source segment' (15)). Such metaphors could be understood as portraying translation as an undifferentiated mass, in which there is little scope to discern the particularities of individual translations, or traces of the work of individual translators (see Kenny 2008, (16)). But at least the food and crop metaphors imply that a human agent has prepared the food, or planted the crop in question. A further metaphor, that of 'translation as a natural resource' (at least one major multinational corporation has described its translation memory assets in this way) leaves no such room for an

understood agent, responsible for the creation of translations. Rather translation is understood as occurring in nature, and like other natural resources, open to exploitation by those with the appropriate technology and legal rights.

As to the sustainability of the translation profession, we again get mixed messages from academia. If SMT cannot survive without a supply of human translations as Way and Hearne suggest (see also Ozdowska and Way 2009, (17)), then the future of human translation does not seem in doubt. But other commentators are less optimistic. Ignacio Garcia (18), who teaches translation at the University of Western Sydney, has claimed that:

as soon as 2010, translation for localisation will be pushed into simple MT post-editing, while other sectors will see a shift toward call-centre conditions and a return of the amateur. (Garcia 2009:211)

While the above prediction is limited to ‘translation for localisation’, in a subsequent article Garcia 2010 (19) claims that students at his own University produced better (English-Chinese and Chinese-English) translations when they post-edited SMT output than when they translated from scratch. This finding is then used to motivate the question of whether translators should consider post-editing as a viable alternative to conventional translation. But one might easily ask about the ethics, from the trainee translator’s point of view, of opting for a machine translation solution over a poorer human effort, rather than endeavouring to improve the human effort in the first instance. Garcia’s response might be an ethical one, if one believes that students are better served in the long run by acquiring post-editing skills rather than translation skills (and remember that Garcia believes that post-editing will be required more than translation even in the short term, in the localization industry at least). But if post-editing is often associated with mere ‘good enough’ quality, what becomes of the ethical translator’s desire to strive for excellence (following Chesterman)? Do we instead have to consider what kind of post-editor an ethical post-editor would be? And is it possible to seek to excel at producing ‘good enough’ post-edited text?

These, and concerns about payment and the sustainability of post-editing as a profession, are certainly questions that have already begun to trouble some translation teachers. But again, opinions differ greatly: Gouadec 2007:25 (20) sees post-editing as ‘a very attractive proposition’ for some translators, while Sharon O’Brien (personal communication) has reported that ‘better’ students found post-editing somewhat tedious in experiments she conducted at Dublin City University. Nonetheless, O’Brien (21) remains an advocate of post-editing in scenarios where appropriate system customization goes a long way to ensuring that post-editors have good SMT output to work with in the first place. She also recommends that post-editing be alternated with regular translation tasks in cases where tedium or fatigue could be an issue. Implicit in O’Brien’s position is the expectation that post-editors are also qualified to work as translators. The assumption that post-editors are at the very least bilingual contrasts starkly to the position of some other researchers in SMT. Philipp Koehn (22), for example, extols the virtues of monolingual target-language post-editors (he actually calls them ‘monolingual translators’). O’Brien’s careful attention to the conditions in which future translators will work serves as an illustration of how translation teachers already

take the interests of current students into account, but one can envisage that ethical considerations will impinge more and more on the teaching of translation over the coming years, if technologization continues at its current pace.

Concluding remarks

It is clear that translation is at something of a crossroads at the moment, with a number of commentators predicting a shift from translation to post-editing. And at a time when machine translation has never been as reliant on human translation as it is now, it is ironic that the role of translators in the creation of parallel data is often obscured. But translators can take some comfort from the fact that they are not the only ones affected by current trends in computing and further afield. As Jaron Lanier (23) has put it, the economics of free content and crowd dynamics mean that authors, journalists, musicians, and artists are also susceptible to having ‘the fruits of their intellects and imaginations’ treated ‘as fragments to be given without pay to the hive mind’ (Lanier 2010:83). How we will deal with these issues is still unclear, but what is clear is that we will need some kind of an ethical basis to help us rise to the challenge.

References

- (1) Kenny, Dorothy and Andy Way, *Teaching Machine Translation & Translation Technology: a Contrastive Study*, In Mikel L. Forcada, Juan Antonio Pérez-Ortiz and Derek Lewis (editors) MT Summit VIII. Proceedings of the Workshop on Teaching Machine Translation, IAMT/EAMT, Santiago de Compostela, pp. 13-17
- (2) Hearne, Mary and Andy Way, *Statistical Machine Translation: A Guide for Linguists and Translators*, Language and Linguistics Compass, 2011, 5, pp. 205-226
- (3) Van der Meer, Jaap, *The emergence of FAUT: Fully Automatic Useful Translation*, Paper presented at the European Association for Machine Translation annual conference, Oslo, 2006. Abstract available from: <http://www.mt-archive.info/EAMT-2006-VanderMeer.pdf>
- (4) Chesterman, Andrew, *Proposal for a Hieronymic Oath*, In Anthony Pym (editor) The Return to Ethics, special issue of The Translator, 2001, 7, 2, pp. 139-154
- (5) Baker, Mona, In Other Words, 2nd edition. London/New York, Routledge, 2011
- (6) MacIntyre, Alasdair, After Virtue. A Study in Moral Theory, Notre Dame, Indiana, University of Notre Dame Press, 1981
- (7) Cheney, George, Daniel J. Lair, Dean Ritz and Brenden E. Kendall, Just a Job? Communication, Ethics & Professional Life, Oxford, Oxford University Press, 2010
- (8) Melby, Alan K. and C. Terry Warner, The possibility of language: a discussion of the nature of language, with implications for human and machine translation, Amsterdam/Philadelphia, John Benjamins Publishing, 1995
- (9) Melby, Alan K. Why Can't a Computer Translate More Like a Person? 1995, Available from <http://www.mt-archive.info/Melby-1995.pdf>
- (10) Marcu, Daniel and Alan K. Melby, *Data-Driven Machine Translation: a conversation with linguistics and translation studies*, paper presented at AMTA 2006: 7th Conference of the Association for Machine Translation in the Americas, “Visions for the Future of Machine Translation”, Cambridge, MA, August 2006, Available from <http://www.mt-archive.info/AMTA-2006-Marcu-Melby.pdf>

- (11) Och, Franz, *Statistical Machine Translation*, Google Tech Talk 30 June 2009. Available from: http://www.youtube.com/watch?v=y_PzPDRPwIA
- (12) Drugan, Jo and Bogdan Babych *Shared Resources, Shared Values? Ethical Implications of Sharing Translation Resources*, In Ventsislav Zhechev (editor) Proceedings of the Second Joint EM+/CNGL Workshop Bringing MT to the User: Research on Integrating MT in the Translation Industry (JEC '10) Denver, CO, 2010, pp. 3-9. Available from: <http://www.mt-archive.info/JEC-2010-Drugan.pdf>
- (13) Rachels, James and Stuart Rachels, The Elements of Moral Philosophy, 6th Edition, New York, McGraw-Hill International Edition, 2010
- (14) Way, Andy and Mary Hearne, *On the Role of Translations in State-of-the-Art Statistical Machine Translation*, Language and Linguistics Compass, 2011, 5, pp. 227-248
- (15) Quah, Chiew Kin, Translation and Technology, Basingstoke and New York, Palgrave MacMillan, 2006
- (16) Kenny, Dorothy, *Instrumental Uses of Parallel Corpora*, paper presented at New trends in corpus linguistics for language teaching and translation studies. In honour of John Sinclair, Granada, Spain, 2008.
- (17) Ozdowska, Sylwia and Andy Way, *Optimal Bilingual Data for French–English PB-SMT*, paper presented at the European Association for Machine Translation annual conference, Barcelona, 2009. Available from: <http://www.mt-archive.info/EAMT-2009-Ozdowska.pdf>
- (18) Garcia, Ignacio, *Beyond Translation Memory: Computers and the Professional Translator*, The Journal of Specialised Translation, 2009, 12, pp. 199-214
- (19) Garcia, Ignacio, *Translating by post-editing: is it the way forward?* Machine Translation, 2011, 25(3), pp. 217- 237
- (20) Gouadec, Daniel, Translation as a Profession, Amsterdam/Philadelphia, John Benjamins, 2007
- (21) O'Brien, Sharon, *Post-editing Tutorial*, Association for Machine Translation in the Americas (AMTA) Conference, Denver, CO, 2010.
- (22) Koehn, Philipp, *Enabling Monolingual Translators: Post-Editing vs. Options*, In Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the ACL, 2010, pp. 537–545. Available from <http://www.mt-archive.info/NAACL-HLT-2010-Koehn.pdf>
- (23) Lanier, Jaron, You are not a gadget: A manifesto, New York, Alfred A. Knopf, 2010