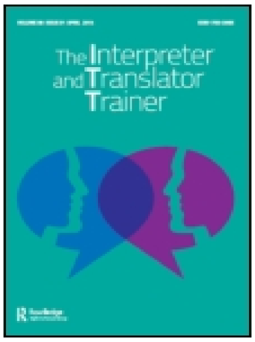

Routledge
Taylor & Francis Group

The Interpreter and Translator Trainer

ISSN: 1750-399X (Print) 1757-0417 (Online) Journal homepage: http://www.tandfonline.com/loi/ritt20

Statistical machine translation in the translation curriculum: overcoming obstacles and empowering translators

Dorothy Kenny & Stephen Doherty

The Interpreter and Translator Trainer, 2014
Vol. 8, No. 2, 276–294, http://dx.doi.org/10.1080/1750399X.2014.936112

Statistical machine translation in the translation curriculum: overcoming obstacles and empowering translators

Dorothy Kenny* and Stephen Doherty

Centre for Translation and Textual Studies, School of Applied Language and Intercultural Studies and Centre for Next Generation Localisation, Dublin City University, Dublin, Ireland

(*Received 11 September 2013; accepted 10 April 2014*)

In this paper we argue that the time is ripe for translator educators to engage with Statistical Machine Translation (SMT) in more profound ways than they have done to date. We explain the basic principles of SMT and reflect on the role of humans in SMT workflows. Against a background of diverging opinions on the latter, we argue for a holistic approach to the integration of SMT into translator training programmes, one that empowers rather than marginalises translators. We discuss potential barriers to the use of SMT by translators generally and in translator training in particular, and propose some solutions to problems thus identified. More specifically, cloud-based services are proposed as a means of overcoming some of the technical and ethical challenges posed by more advanced uses of SMT in the classroom. Ultimately the paper aims to pave the way for the design and implementation of a new translator-oriented SMT syllabus at our own University and elsewhere.

Keywords: Statistical Machine Translation (SMT); human role in SMT; curriculum design; ethics

1. Introduction

The importance of technology in translator education is well established. Several sources argue that there is an onus on translation programmes to help students become confident, flexible and critical users of a variety of computer-aided translation (CAT) tools, so that they will be able to hold their own in an increasingly technologised industry (for fairly recent discussions, see Kenny 2007; Bowker and Marshman 2010; Marshman and Bowker 2012). Initiatives such as the European Master's in Translation (EMT) Network push the point home: translation studies programmes must include substantial training in translation technology (among other things) to be admitted to the Network (EMT Expert Group 2009). But as well as serving this admittedly instrumentalist agenda, it can also be argued that the study of translation tools and their use is valuable because it is inherently interesting: contemporary translation technologies draw on data-driven solutions that reuse existing human translations, raising interesting questions about agency and trust for example (Kenny 2012). They are often used in environments that allow for large-scale collaboration (see contributions to O'Hagan 2011) and play an important role in situated and embodied cognition (Risku 2010). And if translation has always been defined by the technologies it uses (as argued by Cronin 2003, 2012), then it makes little sense to see expertise in the use of new technologies as an add-on that otherwise completely formed student translators can opt into. New technologies mean that translation *itself* changes;

*Corresponding author. Email: dorothy.kenny@dcu.ie

translation now is not what it was 10 or 15 years ago, and this has implications for how we define basic concepts in translation and how we teach translation practices (Pym 2012). This is not to say that all translators are currently using the full range of translation technologies available on the market. Indeed, we would never expect this to be the case, as different technologies (for example, translation memory or machine translation) may be more or less useful depending on a whole host of factors, including the file formats that translators are dealing with, the language pair in question, and the quality levels expected, to name just a few. Rather, the wider field of translation has changed and it behoves translation scholars and teachers to reflect such changes in their scholarship and teaching, while professional translators have to be able to position themselves within changing markets as practitioners who do or do not offer certain translation services (including CAT and MT services), perhaps depending on circumstances.

Such arguments underline *why* we should teach translation technology. An increasing number of publications also present detailed descriptions of *how* particular tools can be incorporated into a more narrowly construed translation technology syllabus, or a more broadly construed translation studies curriculum. Most of the papers in the *Journal of Translation Studies* 2010 special issue on teaching computer-aided translation (Chan 2010) fall into the former category, while work carried out under the banner of the Collection of Electronic Resources in Translation Technologies (CERTT) – at the University of Ottawa takes a broad, holistic view and attempts to create the conditions in which a range of technologies can be easily integrated into courses across the translation studies curriculum (Bowker and Marshman 2010; Marshman and Bowker 2012). The place of machine translation (MT) in most of the pedagogically inspired work referred to here is, however, marginal. Bowker and Marshman (2010, 204) mention, for example, that tutorials and exercises for the teaching of MT, exemplified by the rule-based system (see below) Reverso Pro, have been created as part of the CERTT project, but they do not give any further details. The other papers in Chan (2010) that mention MT say little if nothing about *teaching* MT. Despite Pym's (2012) assertion that 'there has actually been quite a lot of reflection on the ways MT and post-editing can be introduced into teaching practices', one has to go back a decade to find sources that focus specifically on the teaching of MT. Arguably, the heyday of reflection in the area was between 2001 and 2003, when the European Association for Machine Translation (EAMT) and the Association for Machine Translation in the Americas (AMTA) devoted three workshops to the teaching of machine translation (Forcada, Pérez-Ortiz, and Lewis 2001; EAMT/BCS 2002; AMTA 2003). But even the papers presented in these workshops mostly concentrate on technologies other than SMT, the technology that was to revolutionise MT in the decade that ensued, and one with which translation teachers must now contend.[1] The exception here is Kevin Knight's (2003) short paper on resources for introducing concepts of SMT, a contribution that remains valuable to this day.

We return to contemporary approaches to teaching SMT below. For now, we wish to argue that the teaching of MT (and in particular SMT) is an under-researched area, despite the increasing importance of SMT and its inherent interest. Our experience also indicates that SMT is not widely taught on translation programmes;[2] and what literature there is seems to suggest that when SMT *is* incorporated into translator training programmes, the role of 'translators' in the SMT workflow is constructed in a limiting way (see below). In the following we argue for a holistic approach to the teaching of SMT, one that will be more empowering than alternative approaches for those of our students who may embrace the technology in their future careers. We have taken such a holistic approach in our own teaching of SMT at Dublin City University, an experience outlined in brief in Doherty, Kenny, and Way (2012) and in detail in Doherty and Kenny (2014). Below we set the scene for the implementation of a new SMT syllabus by giving a very brief explanation of SMT and by reflecting on the role of humans in SMT workflows. We discuss potential barriers to the use of SMT by translators generally and in translator training in particular, and propose some solutions to the problems thus identified.

2. Statistical machine translation

SMT is the technology behind familiar MT systems such as Google Translate™, Microsoft® Translator and Asia Online™. It is based on an intuitively simple strategy: rather than trying to encode a priori in the form of dictionaries, grammars and knowledge bases, all the linguistic and world knowledge required to translate a text from one language into another (the approach taken in **rule-based** and **knowledge-based** MT), simply *learn* how to translate from already existing human translations. In practice, such learning involves the induction of statistical models of translation from **parallel corpora**, that is, source texts and their human translations. In the terminology of SMT, we say that **translation models** are **trained** on these parallel corpora. SMT systems also rely on so-called **language models**, or monolingual models of the target language; so rather than just ask whether ‘the house’ is a likely translation of *la maison* (the answer to which question should come from the translation model), the SMT system also needs to ask whether ‘the house’ is a likely sequence in English in the first place. Language models can be trained on the target language side of a parallel corpus, or on larger monolingual corpora of target language text.

Following Hearne and Way (2011), we expand here on language models first, as they allow us to introduce important basic concepts in SMT in a relatively straightforward manner.[3] We then reuse some of these concepts in our brief explanation of translation models. Our overview of the technical details of SMT will be necessarily short. The interested reader is referred to Hearne and Way (2011), a paper that was written especially for linguists and translators. Another excellent source is Philipp Koehn’s (2010) textbook *Statistical Machine Translation*. Although it is aimed at computer scientists, much of it is accessible to non-computationally trained linguists.

2.1 Language models

Given a corpus of target language texts, it is possible to create a model of the target language (or, more exactly, a model of the corpus) based on the distribution of single words in the corpus. Such a model is called a **unigram** model. By way of illustration, if we had a tiny corpus consisting of just the single 10-word sentence in (1), we could induce the unigram model in Table 1. Here we make simple observations such as: ‘she’ occurs once in 10 words, while ‘in’, ‘the’ and ‘biggest’ each occur twice in 10 words. We move from these observations of frequency to statements of probability and say (in Table 1) that the probability of ‘she’ occurring is one in ten (or 0.1), the probability of ‘in’ is two in ten (or 0.2), and so on. These probabilities can later be applied to as yet unseen strings (see (2) below for an example).

(1) She lives in the biggest house in the biggest village.

Given the model captured in Table 1, we can compute a probability for whole sentences by simply multiplying their unigram probabilities.[4] The probability of sentence (2)

Table 1. A simple unigram language model..

E	*p*(*e*)
She	0.1
lives	0.1
in	0.2
the	0.2
biggest	0.2
house	0.1
village	0.1

Note: Here *e* stands for an English word, and *p*(*e*) for the probability of that word.

(2) She lives in the village.

is thus:

P(she).P(lives).P(in).P(the).P(village)[5]
= (0.1) × (0.1) × (0.2) × (0.2) × (.01)
= 0.00004 (or 4 in 100,000)

While the probability of:

(3) She lives.

is: P(she).P(lives)
= (0.1) × (0.1)
= 0.01 (or 1 in 100)

According to the very limited model in Table 1, sentence (3) is thus a far likelier sentence of English than sentence (2).

Unigram models are, of course, beset with problems. One is that they systematically give higher scores to shorter sentences. Another is that they do not take word order into account, so (4) will be assigned exactly the same probability as (2):

(4) village the in lives she.

A further problem is that if a previously unseen sentence contains a word that was not present in the training data, then that word is assigned a probability of zero. Take sentence (5), for example:

(5) She lives in the suburbs.

Here, if the unigram probability of 'suburbs' is zero (because it did not appear in the training data on the basis of which we induced the model in Table 1), then the probability of sentence (5) itself is zero. Of course, if we had better (that is, much more) training data, we would have better unigram models, and far fewer words in previously unseen sentences would be deemed to be unknown. But no training corpus will ever contain every word in a given language, so another solution is required. The solution normally adopted is to assign tiny probabilities to

Table 2. Bigrams in sentence (1).

She lives
lives in
in the
the biggest
biggest house
house in
in the
the biggest
biggest village

unknown words such as 'suburbs' in sentence (5), so that sentences containing such words are assigned low probabilities but are not deemed impossible. In other words, a small amount of probability mass is reserved for unseen events.

Bigram models, which are based on sequences of *two* consecutive words, are somewhat more sophisticated. The bigrams present in our tiny corpus (sentence (1) above 'She lives in the biggest house in the biggest village') are reproduced in Table 2.

While 'in the' and 'the biggest' occur twice, all other bigrams occur only once. As before, we can use counts from our corpus to estimate the probability of future events. In this case we calculate the probability of a bigram by dividing its frequency in our corpus by the frequency in our corpus of the first word in the bigram. The probability of the bigram 'she lives' is thus 1/1 = 1; while the probability of 'in the' is 2/2 = 1; and the probability of 'biggest village' is ½ = 0.5. This means that (given our tiny training corpus!) if we see 'she' then we fully expect to see 'lives'. If we see 'in' we fully expect to see 'the'. But if we see 'biggest' there's only a one-in-two chance that the next word will be 'village', as it could equally be 'house'.

The bigram probability of the previously unseen sentence:

(6) She lives in the biggest village.

is again calculated by multiplying the probabilities of its individual bigrams.

P(she lives).P(lives in).P(in the).P(the biggest).P(biggest village)[6]
= 1 × 1 × 1 × 1 × 0.5
= 0.5

This type of modelling can be extended to strings of three words or **trigrams** (see Table 3).

Table 3. Trigrams in sentence (1).

She lives in
lives in the
in the biggest
the biggest house
biggest house in
house in the
in the biggest
the biggest village

Using trigrams we attempt to capture information about how likely we are to see a particular word given the *two* preceding words. We calculate the probability of a trigram by dividing its frequency by the frequency of the first two words in the trigram. The probability of the trigram 'the biggest house' is thus the frequency of 'the biggest house' divided by the frequency of 'the biggest':

$$\frac{\text{frequency ('the biggest house')}}{\text{frequency ('the biggest')}} = \frac{1}{2} = 0.5$$

The trigram probability of a previously unseen sentence is, again, calculated by multiplying the probabilities of its constituent trigrams.

So the trigram probability of sentence (6) above, 'She lives in the biggest village' is

$$P(\text{she lives in}).P(\text{lives in the}).P(\text{in the biggest}).P(\text{the biggest village}) =$$
$$(1) \times (1) \times (1) \times (1) \times (0.5) = 0.5$$

(which happens to be the same as its bigram probability, given our language model).

Language models can also be based on 4-grams, 5-grams, and so on; an upper limit of around seven words is usually imposed on ***n*-grams**, however, due to limitations on computational resources and time constraints. Longer *n*-grams are intuitively better as they are more context-sensitive, but the longer the *n*-grams identified in previously unseen sentences, the less likely it will be that they appeared in the training data. To combine the strengths of shorter, more flexible *n*-grams with longer, more context-sensitive *n*-grams, it is normal to use several language models (unigram, bigram, trigram, etc.) at the same time when computing the probability of a previously unseen sentence, with greater weight usually assigned to longer *n*-grams.

What is most important for the interested translation scholar to note is, however, that in SMT the probability of a target language sentence is calculated based on the joint probabilities of the *n*-grams it contains. This type of modelling is called **generative modelling** and involves breaking a bigger problem up into a series of smaller steps (Koehn 2010, 86). The *n*-grams themselves are not linguistically motivated, that is, they do not correspond to categories in linguistic theory like 'constituent' or even 'collocation'. Sometimes *n*-grams are coterminous with constituents (the trigram 'the biggest house' coincides nicely with the constituent 'noun phase' or 'nominal group', for example) but this is just a coincidence: other trigrams in our training corpus in sentence (1) above ('lives in the' or 'house in the'; see Table 3) are not coterminous with constituents in syntactic theory. A good *n*-gram model might be expected to capture some constituents and collocations,[7] but it will also capture many *n*-grams that have no status in linguistic theory.

2.2 Translation models

Translation models, the second part in our SMT jigsaw, also rely on *n*-grams. They are captured principally in tables in which sequences of source language words are paired with sequences of target language words, and a probability is assigned to each pairing. In unigram models, as should by now be clear, these tables are based on single words. Thus the French word *maison* might be paired with the English word 'house' and a probability assigned to the pairing. For example, a probability of 0.6 would indicate that in six times

Table 4. Translation table for *maison*.

maison	
e	$t(e\|f)$
house	0.6
home	0.3
company	0.1

out of ten, 'house' is found as a translation of *maison* in the parallel corpus that is used to train the SMT system in question. Other translations of *maison* might also be found in the training data, and they too would be assigned probabilities on the basis of their observed frequency. Table 4 represents a very simple **translation table** or **T-table** showing probabilities for three translations of *maison*, where e stands for an English word and $t(e|f)$ denotes the translation probability of the English word e, given the French word f.[8]

Table 4 in some ways suggests that we know all there is to know about how to translate the word *maison*, based on what we have already observed in our (fictional) training corpus. We could (again by reserving some probability mass) allow for the possibility that there are other translations of *maison*, ones that have not been seen yet because our parallel corpus is too small perhaps, or because it does not cover enough different domains, but this need not detain us here.

Table 4 also masks the considerable work that goes into working out which target language words pair with which source language words in the first place. While it might be obvious to a bilingual human that 'home' is the translation of *maison* in examples (7a) and (7b) below, in SMT this usually has to be *learned* directly from the training data by machines in a process known as **word alignment**.

(7a) *La maison familiale doit parfois se vendre à perte.*
(7b) The family home may have to be sold at a loss.
(from www.linguee.com and attributed to erc-cee.gc.ca)

The machine learning algorithm used to produce alignments between words in paired sentences[9] typically starts off by hypothesising that any unigram in (7a) could be paired with any unigram in (7b). So a probability would be assigned to the alignment of: *maison* with 'the'; *maison* with 'family'; *maison* with 'home'; *maison* with 'may', and so on. Successive iterations of the algorithm then suggest that some alignments are much more probable than others.[10] Ideally, the algorithm would ultimately conclude that the alignment of *maison* with 'home' is the most probable. Evidence from other sentence pairs might suggest the alignment of *maison* with 'house', ultimately contributing to the kind of translation table depicted in Table 4. Once word alignments have been worked out, it is possible to work out alignments between larger units, for example, between *maison familiale* and 'family home', leading to what are known in SMT as **phrase alignments**.

Table 5. Sample phrase alignment.

English	↔	French
could not	↔	*impossible*

Table 6. Sample phrase table.

maison familiale	
e	*p*(*e*|*f*)
family home	0.5
family residence	0.3
Home	0.15
their domicile in	0.05

Note: Here *p*(*e*|*f*) stands for the probability (*p*) of an English *n*-gram (*e*), given a French *n*-gram (*f*), *f* being *maison familiale* in this case.

Note, however, that these phrase alignments may be between *n*-grams that would not be considered 'phrases' according to the usual understanding of this term in linguistics. Hearne and Way (2011, 218), for example, give the alignment in Table 5 as an illustration of a phrase alignment.

Phrase alignments, along with their probabilities, are captured in **phrase tables**. An invented phrase table for the bigram *maison familiale* is presented in Table 6.

Phrase tables, containing *n*-grams of varying length and which can record many-to-many correspondences, function as the translation models on which SMT systems draw. As already indicated, the statistical models outlined above are extracted during the **training** phase in SMT. In a subsequent phase, the SMT systems also learn the optimal weighting that should be given to the translation model on the one hand, and the language model on the other, in order for the system to produce the best possible outputs. This process is known as **tuning**. For a more detailed description, see Hearne and Way (2011, 218–219).

Once an SMT system has been trained and tuned, it is ready to start translating new, previously unseen source-language sentences. To do so, the system matches *n*-grams in those sentences against *n*-grams in its phrase table (where these are available), and retrieves their probable translations. *n*-gram translations can be strung together in any order to temporarily produce a number of hypothetical target-language sentences, some of which may be ungrammatical and even nonsensical. Once the (target) language model is taken into account however, some (hopefully all!) of these nonsensical translations are weeded out, because they are highly improbable according to the language model. SMT is thus a matter of generating multiple translation candidates for a given source-language sentence, and then finding the one that is most probable according to both the translation and the target language model. This 'translation proper' phase of SMT is known, somewhat misleadingly (Kenny 2011), as **decoding**.

3. What human translators need to know

The examples used above are extremely simplistic: contemporary Phrase-Based SMT systems consider far more factors (known in SMT circles as **features**) than merely what is predicted by *n*-gram-based translation and language models. It is important for our purposes, however, to understand a number of basic points to do with the role of humans in SMT workflows, and with how we conceive of SMT in translator training. Most of these points follow on from the technical implementation of SMT systems described above.

Firstly, the translation models used in SMT are induced on the basis of parallel corpora of source texts and their *human* translations. Translators (and translation companies) produce such corpora; they are thus already present at the very start of the SMT chain.

As Way and Hearne (2011, 238) have argued: 'the role of the translator in SMT is a crucial one: they provide all the knowledge upon which our models are based'.[11]

Secondly, translation and language models reflect the data on which they were trained. A system trained on legal texts, for example, will be more useful for translating legal texts than it will be for texts from other domains. Translators and translation companies who specialise in particular areas are the very people who are likely to possess the **in-domain** texts on the basis of which useful translation models can be trained for their own needs. Such specialised data constitute a valuable commodity, especially when their quality can be vouched for.

Thirdly, as has already been pointed out, the probabilistic translation and language models used in SMT are *n*-gram based, that is, they rely on word strings of limited length that do not systematically correspond to any type of syntactic constituent. So just as 'probabilistic translation' is a valid bigram in the previous sentence, so too are 'translation and' and 'length that'. Understanding this is crucial to understanding what kind of interventions might be fruitful in any attempts to improve the output of SMT systems. Rather than systematically attempting to resolve syntactic ambiguities in source texts before sending them for MT (which is the kind of thing one might have done with a rule-based MT system), for example, an SMT user might prioritise the creation of subject-specific glossaries that will ensure reliable translation of terms.

Fourthly, and related to the last point, the probabilistic processing steps that underlie the generation of target language sentences in SMT do not have an obvious counterpart in human translation. David Bellos (2011, 266) might be right when he says that professional translators behave like Google Translate™ in 'scanning their own memories in double-quick time for the most probable solution to the issue at hand', and the general point that much language use is *re-use* is well made. But the devil is in the detail, and when one looks at how SMT builds candidate translation sentences, the differences between SMT and human translation are put into sharp relief. As Hearne and Way (2011, 206) put it:

> The idea of generating target sentences by translating words and phrases from the source sentence in a random order using a model containing many nonsensical translations may not seem plausible. In fact, the methods used are not intended (in our opinion, at least) to be either linguistically or cognitively plausible.

This means that 'black-box' evaluations of SMT that critique the technology on the basis of how it deals with given linguistic categories, a particular case or tense in Finnish, for example (see Robinson 2012, 42–44), can be limited in their application. While it is ostensibly true that there are realisations in source texts of linguistic categories such as grammatical role (subject, object, etc.), case (nominative, accusative, etc.), tense (past, present, etc.), and so on, and that the words that realise such categories are 'translated' by SMT systems, 'pure' SMT itself generally has no conception of such categories; and there is no obvious, direct way to intervene in an SMT system in order to change how it deals with case, or tense, and so on. This is not to say that the metalanguage provided by linguistics is not useful when it comes to diagnosing problems in SMT output, rather it is not immediately clear how such diagnoses can feed back into efforts to improve system performance.

Fifthly, as we have seen, translation and language models are sometimes traded off against each other, to strike a balance between 'adequate' translations that actually reflect the source-language input and 'fluent' translations that work well as target-language sentences (because they represent sentences that are very likely to occur in the target language). While a considerable amount of ingenuity goes into tuning systems so that machine translation developers get the balance right between adequacy and fluency (and other features that may be taken into account), users of SMT systems still need to be aware that output that looks good in the target language might not necessarily be a 'faithful' translation of the source input. The only way to be absolutely sure is to have a competent human check. The competent human in question might be a bilingual reader who checks both the source and the target texts, but other types of human evaluation are possible (see, for example, White 2003; Bowker and Ehgoetz 2007).

Such human evaluations are expensive to conduct and inevitably subjective (Koehn 2010, 218–220), and developers of MT systems have developed alternative, automatic ways of gauging how their systems are performing. These **automatic evaluation metrics** (AEMs) for the most part rely on 'gold standard' human reference translations, against which machine outputs are compared.[12] AEMs, like the statistical models used in SMT, tend to be based on *n*-grams, and usually count the number of *n*-grams shared by the machine output and the human reference translation (the higher the number of shared *n*-grams the better the machine output is deemed to be). AEMs have well-known shortcomings (Koehn ibid., 228–229), but they are widely used in SMT circles to measure incremental improvements in system performance while systems are still under development, or to compare systems against each other in 'shared-task' evaluations (see, for example, Bojar et al. 2013). They are also used in the tuning stage of the SMT process. So despite their drawbacks, well-known AEMs such as BLEU (BiLingual Evaluation Understudy, Papineni et al. 2002) remain extremely important to the SMT community, and it is helpful for human translators who engage with SMT to know how these AEMs work. Human translators who have this knowledge can, for example, explain why a human evaluation of a stretch of machine translation output might, on occasion, give a very different result to an automatic evaluation. They might also be able to recommend and use the most appropriate evaluation methods in a given scenario.

Finally, even if an SMT system has been trained on high quality, in-domain data, and tuning has been optimal and evaluation thorough, there may still be a need to fix errors (or make other edits) in translations output by the SMT system. Professional human translators are obvious (but not the only) candidates for this **post-editing** work, given their traditional skill sets (O'Brien 2002). For some translators, post-editing no doubt represents 'a very attractive proposition' (Gouadec 2007, 25); others, however, may view it as reductive (Koehn 2010, 23).

While this section has been necessarily short, what we hope to have emphasised are the stages at which human translators might usefully intervene in the SMT process, and the kinds of knowledge they need to possess in order to do so. This is consistent with what we hope is a holistic, empowering approach to teaching SMT, one that does not exclude human translators from any part of the process in which they could conceivably participate. If translators have the necessary data, know-how and technology, then we see no reason to exclude them from SMT workflows. As we will see below, some sources seek to marginalise translators in SMT workflows, while others allow for only a limited, after-the-event role for translators (typically associated with post-editing machine translation output) in those same workflows. We thus occupy a position that opposes some voices or is only partly consistent with others. In the following sections we attempt to work through some of the tensions in the literature. Before doing so, however, it behoves us to consider how much of a preoccupation SMT should be for translation trainers. Given a history of hype from MT developers and scepticism on the side of many translators,[13] it is worth asking to what extent SMT is used in professional translation workflows, and thus how burning an issue it really is for translation graduates who are about to enter the profession.

4. The use of SMT in professional translation workflows

It is difficult to find reliable, independent statistics on how many translators use MT in general or SMT in particular. That said, we can form partial pictures of the professional use of MT from a number of surveys. A recent Europe-wide survey of over 700 employers of translators, predominantly from the private sector (Optimale 2011), found that only 6% of respondents considered the 'ability to parameter machine translation systems' to be an essential skill in new recruits, while 22% considered it 'important'. This means that 72% of respondents replied that this skill was 'not so important' (35%) or 'not required' (37%). Similarly, Lafeber's (2012) survey of 153 international organisations that employ translators also found that new recruits were not required to have expertise in MT or even in translation memory technology. But both the Optimale and the Lafeber surveys offer only a partial picture of the potential use of MT by translators: for one they survey only employers, and given that the vast majority of translators work freelance (see, for example, Olohan 2007), these sources cannot tell us anything about the technologies used by most translators. Furthermore, they elicit information on skills required at recruitment stage, and thus do not attempt to capture what might be desirable skills later on in a staff translator's career. A more recent survey by market research company Common Sense Advisory (CSA) (Kelly, DePalma, and Stewart 2012) also targets employers, in this case language service providers (LSPs) with two or more employees. Their 2012 annual review of the translation, localisation and interpreting services industry, finds a relatively small yet growing demand for post-editing of machine translation output. They report that post-editing of machine translation output accounted for 2.47% of revenues, or US $828.02 million, in the global language services market in the period under review. In contrast, 'traditional' translation remained the mainstay of LSPs, accounting for 45.70% of revenues (Kelly, DePalma, and Stewart 2012). But while revenues from MT post-editing were small in relative terms, the service constitutes a growing market, and 38.63% of the 1119 respondents to the CSA survey reported that they offered post-editing services. The CSA data thus suggest that post-editing cannot be described as a 'niche' service on the basis of the number of LSPs who offer it, and so graduates with post-editing skills should find plenty of LSPs to whom they can offer these skills. But they also show that post-editing services still have only a modest position when it comes to generating revenue, which may give translation graduates and LSPs pause for thought. If the activity generating this revenue happens to be highly profitable, however, then post-editing MT output could be very attractive for those involved. The CSA survey does not report on tasks other than post-editing that are required in SMT workflows and to which we alluded above (creation of parallel corpora, profiling of training data, intervening in the SMT process, evaluation of outputs, etc.). Despite these shortcomings, we think it safe to say that the the CSA survey confirms that there is growing demand for post-editing services, but that it may not be wise for those who are about to graduate to focus on post-editing at the expense of other 'traditional' translation skills. The data from surveys by Optimale (2011) and Lafeber (2012) also support this position.[14]

5. SMT in translator training

We have argued above for an approach to the teaching of SMT that does not exclude human translators from any part of the process in which they could conceivably participate. We feel compelled to promote such a view in the face of what could be described as mild antagonism towards translators from some business-focused SMT quarters. In a post to the Automated Language Translation Group on Linked In, for example, Dion Wiggins (2011), the CEO of the SMT provider Asia Online™, argues for a limited role for translators in translation workflows, one in which the translators would not have any ownership, never mind control:

> The translator should not have any ownership in the translation process. They are 1 part of the translation process. There is much more than the translator. Ownership should be with the level of the LSP and the client, not at the translator level ... Allowing the translator (who is usually a freelancer and also works for your competitors) to control the translation process is a recipe for disaster.[15]

As already suggested, the small but growing body of literature on SMT and translator training, while not antagonistic towards translators in the way that Wiggins is, can nonetheless also construct the translator's role in SMT workflows in a somewhat limited way. This is due to the tendency to see post-editing as the only role for translators (who then morph into post-editors) in such workflows. Recent work by Ignacio Garcia (2010), for example, focuses on differences in performance when a group of 14 students translate texts from scratch as opposed to when they post-edit SMT output. He concludes rather tentatively that 'The data point to the possibility that translators may achieve higher quality when working from the MT baseline' (17), a conclusion that allows him to speculate that 'the question, in the long term, will not be *whether* translation will be done from the MT baseline, but simply *when*' (18). In a related, larger study that considers variables such as direction of translation (into L1 or into L2), text difficulty, and participant type (weaker vs stronger performers), Garcia (2011, 218) goes on to conclude that 'even without any training in post-editing, translation trainees may on average do better when translating via post-editing'. Garcia provides us with useful models for evaluating student performance under the translation and post-editing conditions, but there is a risk that by valorising post-editing and suggesting that 'translation by post-editing' (218) is an inevitability, Garcia may encourage us to forget about the other stages in the translation workflow to which (trainee) translators could easily contribute, given their particular skills. Given the continued prevalence in the market of 'traditional' translation (as shown above), it would also seem unwise not to focus considerable energy on ensuring that trainee translators can deliver high quality even without the use of SMT. Nor does Garcia say anything about *teaching* SMT in the broader sense (as envisioned in this paper), and his comments on teaching post-editing are limited to the observation that '[t]he pioneering research of O'Brien (2002), Belam (2003) and Depraetere (2010) ... needs to be expanded to gain information on the skill set that best suits the task and the ways to develop such skills' (Garcia 2011, 229).

Of the three sources mentioned by Garcia, O'Brien's (2002) treatment of the teaching of post-editing, although not guided by classroom experience, is the most expansive.[16] It includes a justification for teaching post-editing and consideration of who might make good post-editors (translators or others) and what skill sets they might need. It concludes with a proposed syllabus for a post-editing course. This syllabus, however, covers much more than post-editing. It also takes in MT in general, controlled language authoring, advanced terminology management, text linguistics (a term that may date the paper somewhat) and basic programming. It thus covers the entire MT workflow, and is more consistent with the kind of approach to teaching SMT that we envisage here. O'Brien (2002) does not mention SMT per se, however, as this source predates the widespread availability of SMT solutions.

Given the considerable potential for trainee and professional translators to be involved in nearly all stages of the SMT translation workflow, along with interesting differences in attitudes to translator involvement in these workflows, and obvious gaps in the pedagogically oriented literature, we believe that the time is ripe to develop and publish an up-to-date holistic syllabus in SMT for trainee translators. We have developed and implemented such a syllabus at Dublin City University and report in detail on this project in Doherty and Kenny (2014). In the rest of this paper we consider the barriers we had to overcome in implementing this syllabus, and that others translator trainers would have to overcome to do likewise. We also propose some ways to overcome these barriers.

6. (Overcoming) barriers to the use of SMT in translator training

We hope that we have provided sufficient motivation for introducing a holistic syllabus in SMT to the translation curriculum. For many translator trainers and translation students (not to mention professional translators), however, there are significant barriers to the adoption of SMT. Some of these barriers are conceptual: SMT can be difficult to understand, even for computational linguists (see Way 2009; Hearne and Way 2011). Careful, repeated reading of descriptions of SMT intended for linguists and translators (see above) can help trainee translators to overcome this barrier, however, as can attending classes on SMT in particular or probabilistic natural language processing in general.[17] While it is true that users do not need much or even any understanding of SMT to use systems like Google Translate™ (Robinson 2012, 39), Google Translate™ was not initially designed for professional translators,[18] and there are several reasons why professional translators might want to avoid translating anything but short fragments of texts using this source. They may not be able to ensure client confidentiality, for example, and non-disclosure agreements may explicitly prohibit translators from uploading client material to web-based translation services.[19] They may also have qualms about reusing the work of other unacknowledged translators, who may not have given permission for their translations to be used in this manner. Such issues constitute legal and ethical barriers to using (free web-based) SMT in a professional capacity. They are discussed in more detail by Drugan and Babych (2010). One way of overcoming these barriers is not to rely on web-based services, but to build one's own in-house SMT system, but this is where translators may encounter technical barriers to SMT use. Such barriers are expanded on below. For now let us note that Google Translate™ may be paradoxically *too easy* to use; not only does free online machine translation imply that translation is 'an agentless, automatic function that can be realized in no time at all' (Cronin 2012, 47) thus obscuring the human labour that produces the translated and other data on which SMT is based; systems like Google Translate also obscure the labour of the computer scientists who build SMT systems, and can give the impression that there is nothing to SMT.

But if Google Translate™ is too easy to use, other solutions appear to be too hard: SMT was originally the domain of large corporations (Google, Microsoft, etc.) or specialized companies (e.g. Language Weaver, acquired by SDL in 2010) who had at their disposal huge computing power, in-house teams of computer scientists, and either massive corpora of existing translation data or the ability to mine the web for such data. Increasing computer power and the sharing of ideas, data and tools, however, soon allowed the academic community working on SMT to flourish; particularly important was the advent of the open-source Moses toolkit,[20] described by its creators as:

> a complete out-of-the-box translation system for academic research. It consists of all the components needed to preprocess data, train the language models and the translation models. It also contains tools for tuning these models … and evaluating the resulting translations using the BLEU score … (Koehn et al. 2007, 178)

But even Moses, which was specifically designed with accessibility and ease of use in mind (Koehn et al. 2007, 178), was still intended for computer scientists, and even then some well-qualified users found it challenging to implement. In March 2009, Tom Hoar (Hoar 2009, 2) wrote in a Translation Automation Users Society (TAUS) report:

> Let's face it, SMT systems are complex. This is a day when a novice computer user can install and set up a complex office application in 30 minutes. Yet, one experienced *C*++ software engineer actively developing open source projects for Ubuntu Linux took three days to grasp the concepts, collect the dependencies, compile the components, and verify the baseline Moses Decoder application before he could set up an existing trained data model and test the quality of the translation.

Despite ongoing development and detailed documentation, Moses can still be intimidating for non-computer scientists, and it certainly cannot be assumed that translation students with little background in computing will have the knowledge or skills required (or the time required to attain such knowledge and skills, given packed translation studies curricula) to successfully install and run the software.

Not surprisingly perhaps, the launch of Moses was followed in due course by that of other projects, intended to make the technology more accessible. These projects include *Do Moses Yourself*™ and *Moses for Mere Mortals*. An early positive report on the use of the latter system was published by Machado and Fontes (2011), who – importantly for our purposes – evaluated it 'from a translator's perspective' (2). But while implementing *Moses for Mere Mortals* brought clear productivity gains to the Portuguese translation team at the European Commission's Directorate General for Translation where Machado and Fontes work, it is clear that even installing the system remains challenging for those with little knowledge of the Linux™ operating system or scripting languages. Similar issues arise with *Do Moses Yourself*™. A very motivated professional translator may be able to overcome these problems, but the knowledge required to get started with what has become known as 'Do It Yourself' or 'DIY' SMT cannot (yet) be assumed among the cohorts of trainee translators we typically teach. Nor can it necessarily be assumed among translation teachers. In the long term, translation students and their teachers may need to acquire greater computing skills, but in the short term other, cloud-based solutions are possible. Such cloud-based solutions turn out to be 'just right' for our purposes, and in a nice illustration of the Goldilocks principle, we gravitate towards them in our teaching.

7. Cloud-based SMT

The term 'cloud computing' describes the delivery of computing services over a network, typically the Internet. Recently launched cloud-based SMT services include: KantanMT™, developed by Xcelerator Machine Translation Solutions Ltd; Microsoft® Translator Hub; SDL's BeGlobal™; and Capita's SmartMATE™.[21] These cloud-based services allow users to use their own bilingual and monolingual data, sometimes on top of the service provider's 'stock' data, to train customised SMT engines. Because the engine itself remains in the cloud, there is no need to install software locally. Nor is there any need for users to have particularly powerful computers, as all the 'heavy lifting' is done by remote machines, in remote centres such as those run by Amazon Web Services.[22] While users stay at one remove from the SMT engine, compared with DIY solutions, they can still go through the whole cycle of: uploading data; training and testing the engine; where necessary making interventions aimed at improving quality (e.g. through adding more training data, or deploying project-specific glossaries, etc.); retraining the engine; and finally deploying the engine 'for real'. In general, user interfaces are easy to use and allow collaboration between reviewers/testers. In some cases, custom-built translation engines can be kept private, shared with other named parties, or made available publicly.

On the face of it, SMT in the cloud offers a number of advantages for University-based translator training: it allows the user – who could be a translator, translation company, or trainee translator with access to a considerable amount of training data – to intervene at all relevant stages of the SMT workflow, affording ample opportunity for the kind of holistic learning in which we are interested. To our knowledge, however, the use of cloud-based SMT had not been tested in a university translator-training environment until we deployed the service on the MA in Translation Studies and MSc in Translation Technology at Dublin City University in early 2012. In a companion paper to this one (Doherty and Kenny, 2014) we provide more detail on how we designed, implemented and evaluated our SMT syllabus.

8. Conclusions

In this paper we have provided a non-technical overview of SMT, with a view to highlighting the knowledge and skills that trainee translators in particular need in order to work with this technology. We have taken as axiomatic that translators are de facto involved, or can contribute, at nearly all points in the SMT workflow, and that once certain barriers are overcome, there is no need either to exclude translators or to assign them to what some perceive to be limited or reductive roles in SMT workflows. These are crucial points for us: as academics who train translators we have a vested interest in, but also an ethical commitment to ensuring the sustainability of the profession. Otherwise we are educating students in our own interest but with little regard for theirs. (We could, of course, stop educating translators altogether, but this would make little sense given our assessment of the current translation market presented above, among other things.) Such sustainability must partly reside in translators' ability to evolve and to adopt whatever tools are useful and currently available to them so as to remain relevant and competitive. At the same time the adoption of such tools should not force translators into roles that they may find so limiting as to no longer be interesting or professionally rewarding. This is why we advocate a pro-active, holistic approach to the teaching and learning of SMT, one in which translators have ownership, critical understanding and a good deal of control. This, of course, requires those same translators—and those who train them—to acquire a concomitant degree of specialised knowledge and skill, and to overcome barriers, be they conceptual, ethical or technical. We hope to have provided sufficient information and motivation in this paper to encourage other educators to embrace SMT in a similarly holistic way. In a companion paper (Doherty and Kenny, 2014), we show how many of the ideas expounded here can be operationalised, and how a full SMT cycle, from training to evaluation and improvement, can be implemented in the translation curriculum. For now, we hope to have made the case for a reversal of 'disintermediation' in SMT (Cronin 2012, 45–47), in other words, for the reinstatement of the translator as an agent who is very much present throughout the SMT workflow, and for an SMT syllabus that reflects this.

Acknowledgements

This research is supported by the Science Foundation Ireland (Grant 07/CE/I1142) as part of the Centre for Next Generation Localisation (www.cngl.ie) at Dublin City University. The authors also wish to thank our two anonymous reviewers for their invaluable feedback.

Notes

1. It must also be acknowledged, however, that despite its widespread adoption, SMT does not offer a single solution for all MT challenges and much research is currently focused on combining the strengths of earlier rule-based systems with newer SMT systems, in 'hybrid' and 'multi-engine' systems.
2. This impression has been formed partly through extensive contact with translator training programmes that are members of the EMT Network and the Optimale Network.
3. Indeed, the presentation of SMT here owes much to that of Hearne and Way (2011).
4. We *multiply* unigram probabilities based on the assumption that unigrams are *independent* of each other. In probability theory the joint probability of two or more independent events is simply the product of their individual probabilities. Of course this assumption can be easily criticised in the case of natural language text, as words are likely to exert some sort of influence on each other, so 'the book' is more likely than 'the this' for example. Very local influences are better captured in bigram models, addressed below.
5. Here 'P' stands for the probability of the associated word.
6. Note that bigram probabilities are normally represented using the notation P(B|A), which indicates the probability of the second word given that the first word has already occurred. So the probability of the bigram 'she lives' would be expressed as follows: P(lives|she).
7. Here we rely on Manning and Schütze's (1999, 151) definition of a collocation as 'an expression consisting of two or more words that correspond to some conventional way of saying things'. An oft-quoted example is the expression 'strong tea', which is conventionally used in English rather than other potential expressions, such as 'powerful tea'.
8. We follow the notation used in Koehn (2010).
9. A sentence-aligned parallel corpus is required for word alignment.
10. See Hearne and Way (2011) for a worked example of what is known as the **expectation-maximisation algorithm**.
11. It is, of course, possible to train a translation model on a corpus of source texts and their machine translations (see Callison-Burch and Osborne 2003), but some sources have expressed concern about the danger of training translation models on machine translated output that has not been sufficiently 'cleaned' or 'proofed', in other words about the danger of using 'dirty data' (see comments by Dion Wiggins at https://www.taus.net/content/the-new-smt-player-comes-from-thailand).
12. An alternative approach is to try to predict, using 'quality estimation' techniques (see Specia, Raj, and Turchi 2010), how an SMT system will perform on a given occasion given what is known about the system itself, the language pair involved, etc.
13. Such hype is evidenced in one major LSP's assertion (at the Localisation Research Centre Conference in Limerick, Ireland, in September 2010) that contemporary technologies could mean 'no new translation' for their clients. Translator scepticism is often evident in interventions at such conferences or in contributions to forums such as ProZ.com or groups on Linked In, etc.
14. Later surveys provide additional support for an increase in the uptake of MT and the growth of post-editing across the translation and localisation industry. Almaghout et al. (2012) find that, of 438 translators and translation providers, 42% currently use MT. For those who do not use MT, 40% expected to implement it within the next couple of years, and 42% in the longer term. In addition, in a survey of 467 translation/localisation buyers and vendors, Doherty et al. (2013) find that 34% are currently using MT, and 35% plan to employ it in the future: 13% within the next year and 22% thereafter. The survey also finds that 77% of participants state there is an increased need for them to provide MT to their clients.
15. In the same online conversation, Andy Way welcomes the prospect of new 'pay as you go' business models opening up SMT to individual translators. From the point of view of market development, Way's view makes sense: if small LSPs and individual translators can

themselves be consumers of SMT then the market for SMT services becomes much larger than would otherwise be the case, and this can only be good for SMT developers and those selling SMT services.

16. Belam (2003) and Depraetere (2010) are concerned with how post-editing rules or recommendations can be co-written by students and teachers, or else induced by observing the post-editing behaviour of students.
17. Probabilistic techniques are used in a variety of processes relevant to translators, including voice recognition, terminology extraction and SMT.
18. Perhaps unsurprisingly, Google Translate™ is used by professional translators who work for Google, through the Google Translator Toolkit™ interface (Jana Vořechovská, Google, personal communication).
19. It must be said, however, that Google and Microsoft have both gone to some trouble to allay fears about confidentiality when it comes to using Google Translate™ and Microsoft® Translator, respectively.
20. See http://www.statmt.org/moses/. (All web links in this paper were last accessed in March 2014.)
21. For details see: www.kantanmt.com, https://hub.microsofttranslator.com/, http://www.sdl.com/products/sdl-beglobal/and https://www.smartmate.co/.
22. http://aws.amazon.com/.
23. Here e stands for an English word, and $p(e)$ for the probability of that word.

References

Al Maghout, H., F. Gaspari, D. Groves, and J. van Genabith. 2012. *Report on Requirements Gathered from Relevant Stakeholder Groups,* QTLaunchpad Report. http://www.qt21.eu/launchpad/sites/default/files/QTLP-Deliverable-6_5_1.pdf

AMTA. 2003. *MT Summit IX: Workshop on Teaching Translation Technologies and Tools*, Association for Machine Translation in the Americas. http://www.dlsi.ua.es/t4/docum/proceedings.pdf

Belam, J. 2003. "'Buying up to falling down': a deductive approach to teaching post-editing." In AMTA (2003), 1–10.

Bellos, D. 2011. *Is That a Fish in Your Ear? Translation and the Meaning of Everything*. London: Particular Books.

Bojar, O., C. Buck, C. Callison-Burch, C. Federmann, B. Haddow, P. Koehn, C. Monz, M. Post, R. Soricut, and L. Specia. 2013. "*Findings of the Eighth Workshop on Machine Translation*," August 8–9, Sofia, Bulgaria. http://www.statmt.org/wmt13/WMT-2013.pdf

Bowker, L., and M. Ehgoetz. 2007. "Exploring User Acceptance of Machine Translation Output: A Recipient Evaluation." In *Across Boundaries: International Perspectives on Translation Studies,* edited by Dorothy K. and K. Ryou, 209–224. Newcastle upon Tyne: CSP.

Bowker, L., and E. Marshman. 2010. "Toward a Model of Active and Situated Learning in the Teaching of Computer-Aided Translation: Introducing the CERTT Project." *Journal of Translation Studies, Special Issue on the Teaching of Computer-Aided Translation* 13 (1&2): edited by Sin-Wai Chan, 199–226.

Callison-Burch, C., and M. Osborne. 2003. "Bootstrapping Parallel Corpora." In *HLT-NAACL 2003 Workshop: Building and Using Parallel Texts*, 44–9. http://www.mt-archive.info/HLT-NAACL-2003-Callison-Burch.pdf

Chan, S. -W. ed. 2010. *Journal of Translation Studies*. *Special Issue on The Teaching of Computer-Aided Translation* 13 (1&2).

Cronin, M. 2003. *Translation and Globalization*. London: Routledge.

Cronin, M. 2012. *Translation in the Digital Age*. London: Routledge.

Depraetere, I. 2010. "What counts as useful advice in a university post-editing training context? Report on a case study." *EAMT*, May 2010 St Raphael, France. http://www.mt-archive.info/EAMT-2010-Depraetere-2.pdf

Doherty, S., F. Gaspari, D. Groves, J. van Genabith, L. Specia, A. Burchardt, A. Lommel, and H. Uszkoreit. 2013. *Mapping the Industry I: Findings on Translation Technologies and Quality Assessment*. QT LaunchPad Report. http://www.qt21.eu/launchpad/sites/default/files/QTLP_Survey2i.pdf

Doherty, S., and D. Kenny. 2014. "The Design and Evaluation of a Statistical Machine Translation Syllabus for Translation Students." *The Interpreter and Translator Trainer*. doi:10.1080/1750399X.2014.937571

Doherty, S., D. Kenny, and A. Way. 2012. "Taking Statistical Machine Translation to the Student Translator." In *Proceedings of the Tenth Biennial Conference of the Association for Machine Translation in the Americas,* 1–10. http://www.mt-archive.info/AMTA-2012-Doherty-1.pdf

Drugan, Jo., and B. Babych. 2010. "Shared Resources, Shared Values? Ethical Implications of Sharing Translation Resources." In *Proceedings of the Second Joint EM +/CNGL Workshop: Bringing MT to the User,* edited by V. Zhechev, 3–9. http://www.mt-archive.info/JEC-2010-Drugan.pdf

EAMT/BCS. 2002. *Proceedings of the 6th EAMT Workshop: Teaching Machine Translation*. European Association for Machine Translation and the British Computer Society. http://mt-archive.info/EAMT-2002-TOC.htm

EMT Expert Group. 2009. *Competences for professional translators, experts in multilingual and multimedia communication*. http://ec.europa.eu/dgs/translation/programmes/emt/key_documents/emt_competences_translators_en.pdf

Forcada, M. L., J. A. Pérez-Ortiz, and D. Lewis. Eds. 2001. "*MT Summit VIII. Proceedings of the Workshop on Teaching Machine Translation*," IAMT/EAMT. http://mt-archive.info/MTS-2001-TOC.htm

Garcia, I. 2010. "Is Machine Translation Ready Yet?" *Target* 22 (1): 7–21. doi:10.1075/target.22.1.02gar.

Garcia, I. 2011. "Translating by Post-Editing: Is It the Way Forward?" *Machine Translation* 25 (3): 217–237. doi:10.1007/s10590-011-9115-8.

Gouadec, D. 2007. *Translation as a Profession*. Amsterdam: John Benjamins.

Hearne, M., and A. Way. 2011. "Statistical Machine Translation: A Guide for Linguists and Translators." *Language and Linguistics Compass* 5: 205–226. http://onlinelibrary.wiley.com/doi/10.1111/j.1749-818X.2011.00274.x/pdf. doi:10.1111/j.1749-818X.2011.00274.x.

Hoar, T. 2009. *Technical Guide to SMT Training*. TAUS Report, March 2009.

Kelly, N., D. A. DePalma, and R. G. Stewart. 2012. *The Language Services Market: 2012*. Lowell, MA: Common Sense Advisory.

Kenny, D. 2007. "Translation Memories and Parallel Corpora: Challenges for the Translation Trainer." In *Across Boundaries: International Perspectives on Translation,* edited by D. Kenny and K. Ryou, 192–208. Newcastle-upon-Tyne: Cambridge Scholars Publishing.

Kenny, D. 2011. "A corpus-based study of the discourse of contemporary machine translation." Paper delivered at Corpus Linguistics 2011, Birmingham, UK, 20–22 July 2011.

Kenny, D. 2012. "The Ethics of Machine Translation." In *Proceedings of the NZSTI Annual Conference 2011,* edited by S. Ferner, 121–131. Auckland: NZSTI.

Knight, K. 2003. "Teaching Statistical Machine Translation." In *MT Summit IX: Workshop on Teaching Translation Technologies and Tools*, 17–19. http://www.dlsi.ua.es/t4/docum/proceedings.pdf

Koehn, P. 2010. *Statistical Machine Translation*. Cambridge: CUP.

Koehn, P., H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, et al. 2007. "Moses: Open Source Toolkit for Statistical Machine Translation." In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL), Companion Volume Proceedings of the Demo and Poster Sessions,* 177–180. Prague: ACL.

Lafeber, A. 2012. "Translation Skills and Knowledge – Preliminary Findings of a Survey of Translators and Revisers Working at Inter-governmental Organizations." *Meta* 57 (1): 108–131. doi:10.7202/1012744ar.

Machado, M. J., and H. L. Fontes. 2011. "Avaliação do Sistema de Tradução Automática Moses no Departamento de Língua Portuguesa da DGT Utilizando os scripts. Moses for Mere Mortals—Estudo de caso inglês-português' [Evaluation of the Moses Machine Translation System in the Portuguese Language Unit of the DGT, using scripts. Moses for Mere Mortals – an English-Portuguese Case Study]." *A Folha. Boletim da língua portuguesa nas instituições europeias* 37: 12–22.

Manning, C., and H. Schütze. 1999. *Foundations of Statistical Natural Language Processing.* Cambridge, MA: MIT Press.

Marshman, E., and L. Bowker. 2012. "Translation Technologies as Seen Through the Eyes of Educators and Students: Harmonizing Views with the Help of a Centralized Teaching and

Learning Resource." In *Global Trends in Translator and Interpreter Training,* edited by S. Hubscher-Davison and M. Borodo, 69–95. London: Continuum.

O'Brien, S. 2002. "Teaching Post-Editing: A Proposal for Course Content." In *EAMT/BCS*, 99–106.

O'Hagan, M. Ed. 2011. *Linguistica Antverpiensia, New Series - Themes in Translation Studies: Translation as a Social Activity - community translation 2.0*. Antwerp: Antwerp University, Vol. 10/2011.

Olohan, M. 2007. "Economic Trends and Developments in the Translation Industry. What Relevance for Translator Training?" *The Interpreter and Translator Trainer* 1 (1): 37–63. doi:10.1080/1750399X.2007.10798749.

Optimale. 2011. *Optimale Employer Survey*. Results http://www.translator-training.eu/optimale/attachments/article/36/PPT_WP5_survey%20results.pdf

Papineni, K., S. Roukos, T. Ward, and W. -J. Zhu. 2002. "BLEU: A method for automatic evaluation of machine translation." In *Proceedings of ACL-2002: 40th annual meeting of the association for computational linguistics*, Philadelphia, July 2002, 105–114.

Pym, A. 2012. "Translation skill-sets in a machine-translation age." http://usuaris.tinet.cat/apym/on-line/training/2012_competence_pym.pdf

Risku, H. 2010. "A Cognitive Scientific View on Technical Communication and Translation: Do Embodiment and Situatedness Really Make a Difference?" *Target* 22 (1): 94–111. doi:10.1075/target.22.1.06ris.

Robinson, D. 2012. *Becoming a Translator.* 3rd ed. London: Routledge.

Specia, L., D. Raj, and M. Turchi. 2010. "Machine Translation Evaluation Versus Quality Estimation." *Machine Translation* 24 (1): 39–50. doi:10.1007/s10590-010-9077-2.

Way, A. 2009. "A Critique of Statistical Machine Translation." In *Linguistica Antverpiensia, New Series - Themes in Translation Studies: Evaluation of Translation Technology,* edited by W. Daelemans and V. Hoste, 17–41. Antwerp: Antwerp University.

Way, A., and M. Hearne. 2011. "On the Role of Translations in State-of-the-Art Statistical Machine Translation." *Language and Linguistics Compass* 5: 227–248. doi:10.1111/j.1749-818X.2011.00275.x.

White, J. S. 2003. "How to Evaluate Machine Translation." In *Computers and Translation. A Translator's Guide,* edited by H. Somers, 211–244. Amsterdam: John Benjamins.

Wiggins, D. 2011. blogpost to *Automated Language Translation Group*, 13 October 2011. http://www.linkedin.com/groups/Looks-like-licencebased-model-MT-148593.S.74453505?qid=579815d2-fdfd-46bb-ac04-3530d8808772andtrk=group_search_item_list-0-b-ttl