Explaining real MT to translators: between compositional semantics and word-for-word

Mikel L. Forcada,
Dept. de Llenguatges i Sistemes Informàtics,
Universitat d'Alacant, E-03071 Alacant, Spain.
E-mail: mlf@dlsi.ua.es

Abstract

Real (i.e., working) machine translation may be presented both as the result of inevitable approximations over an ideal, theoretically motivated model based on the *principle of semantic compositionality* and as the result of a set of necessary refinements over a very rudimentary *word-for-word* substitutional system. This paper explores the pedagogical value of presenting real MT as being somewhere in the middle of these two extreme scenarios. I contend that it is possible to reshape the (either optimistic or pessimistic) expectations of students about real MT by showing students, on the one hand, the approximations, compromises, and sacrifices necessary to mechanize efficiently a linguistically accurate model, and, on the other hand, the large amount of work needed to improve a word-for-word model so that it produces reasonable translations. This prepares them to learn and appreciate the strategies used to tackle the problem of machine translation.

1 Introduction

One of the main problems when teaching MT to translators and linguists is their attitude towards this technology, which may be partly due to the disappointment of high expectations promoted by MT manufacturers.

For linguists and translators to be able to understand and value the intrinsic difficulties of MT and the achievements of state-of-the-art systems, it is important to place real MT as being in the middle of two very distinct extreme scenarios, which have to be carefully explained: an *ideal model* based on an accurate application of semantic compositionality and a very rudimentary word-for-word, mechanical, substitutional system (*model zero*).

2 MT strategies as approximations to an ideal model

2.1 The principle of semantic compositionality

Many students think that semantics "studies the meanings of *words*", but know too well that translating is more than translating words. Semantics has, however, a lot to say about translation: the *principle of semantic compositionality* (PSC, Radford et al. 1999,

p. 359) states that the interpretation (meaning) of a sentence is compositionally built from the interpretation of its words, following the groupings dictated by its parse tree.

The PSC formulation suggests that interpretations have a structured nature because they are built by composition of simpler interpretations; it is therefore an *ascending* or *bottom-up* formulation (interpretations are built from leaves to root). The PSC is crucial in linguistics because it explains why humans can assign interpretation to (i.e., understand) sentences they have never heard or read before.

A more general and symmetric account of the PSC (Tellier, 2000), which works not only from sentence to interpretation but also in reverse, says that there exist two *bidirectional* mappings:

1. a mapping between *words* and *elementary interpretations* and
2. a mapping between *syntactic rules* (which indicate how constituents are built from other constituents) and *semantic composition rules* (which indicate how partial or total interpretations are built from the corresponding partial interpretations).

The PSC, in its symmetric formulation, is important in linguistics because it provides an explanation to the fact that natural language syntax is structured (that is, compositional) and even recursive: the reason is that the interpretations that humans need to convey are also compositional and recursive[1]. But the symmetric formulation of the PSC is also useful to describe the translation process, as will soon be made clear.[2]

2.2 Compositional translation

In this view, translating a written sentence from SL to TL consists in producing a TL sentence having the same *interpretation* or meaning. Following this scheme, the translation of a sentence proceeds in two stages: the first one obtains a representation of the interpretation of the SL sentence, and the second builds from this representation the corresponding TL sentence. Both in general linguistics and in computational linguistics this representation is usually called *logical form* (Radford et al., 1999, p. 362; Allen, 1995, cap. 8).

Note that reducing the translation of a text to translating one by one its sentences is already an approximation: one which is common to most machine translation systems. It is indeed a rather radical approximation, which, for instance, ignores the structure of discourse. It should however be made clear that the PSC does not explicitly exclude important phenomena which are usually described as intersentential, such as co-reference processes (anaphora). For example, in the PSC framework, anaphoric pronouns may be

[1] Consider the sequence of sentences *Ann was lying*, *John said that Ann was lying*, *Tom said that John said that Ann was lying*, etc.

[2] Often the translation of certain source-language constructs or even the whole translation process is described as being *non-compositional*. It should be clear that semantic compositionality does not mean one can directly combine the *translations* (that is, target language word strings) of constituents of a sentence to obtain the translation of larger constituents but applies rather to the construction of structured, perhaps language-independent *interpretations*.

seen as polysemic words whose interpretation is determined by co-reference with a noun phrase in the same sentence or somewhere else in the text, or by deixis to an object outside the text (*you*, *I*).

In addition, sentences may contain null constituents such as zero pronouns (for example in Japanese) or deleted constituents (as in coordinative and comparative sentences: *Scotsmen like whisky more than Welshmen [like whisky]*, Radford et al. (1999, p. 400)). The covert (non overt) realization of some constituents is not incompatible with the PCS, but does indeed constitute a challenge in MT system design. For example, to translate correctly from Japanese to English one has to assign antecedents to zero pronouns, since Japanese systematically deletes subjects, objects, etc., from sentences when context allows the listener / reader to determine them (Mori et al., 1999; Nakaiwa, 1999).

2.3 Compositional machine translation

The ideal MT system would parse each SL sentence, look up the interpretations for each word, and apply the PSC recursively by walking up the parse tree to build a structured interpretation of the sentence. Then it would apply the reverse PSC by analysing this structured interpretation to obtain a TL parse tree and TL words, and eventually generate the TL sentence from the TL parse tree and words. This is basically the design of *interlingua* systems. The process of translating one sentence may therefore be described as in figure 1.

SL sentence
↓
[Analysis]
↓
SL parse tree + SL words
↓
[PSC for SL]
↓
logical form
(interpretation)
↓
[reverse PCS for TL]
↓
TL parse tree + TL words
↓
[Generation]
↓
TL sentence

Figure 1: Translating a sentence following the principle of semantic compositionality (PCS).

2.4 The obstacle of ambiguity

In this discussion, I have been talking about "the interpretation" of a sentence, seemingly implying each sentence has a single interpretation. However, one of the most important features of human language is its *ambiguity*, a feature that shows clearly at the sentence level. From the PSC point of view, a sentence may be ambiguous:

- because one or more of its words may have more than one interpretation (*lexical ambiguity*),
- because the sentence has more than one parse tree (*structural* or *syntactic ambiguity*),
- or, in some cases, because of both reasons.

The choice of the correct interpretation of a sentence by the MT system, necessary in many cases to produce an adequate translation, is a task which is all but trivial. While humans can use context and their beliefs about the world to safely discard many interpretations (ideally all but one of them), machine translation systems have to make these decisions using only feasible and programmable rules which process the (usually incomplete) information they are capable to extract from the surrounding text.

2.5 The transfer approximation

According to the account above, the work of a professional translator consists in reading each sentence in the SL text, determining its correct interpretation (discarding inadequate interpretations for ambiguous sentences), and building a TL sentence which represents adequately this interpretation. Therefore, the translator must completely understand each sentence before translating it. This is clearly implausible, because one cannot expect all translators to be experts in mechanics or integrated circuit design to translate documents belonging to these subjects (does a translator really need to understand thoroughly what a *crankshaft* or a *flywheel* are to translate engine manuals into Spanish? Or is it enough for her to know that they are equivalent to *árbol de levas* and *volante de inercia*?). In these cases, translators operate by *transforming structures* or *patterns* and *substituting words* (paying special attention to the specific terminology), using rules and strategies that allow them to produce adequate translations without having to understand texts in depth. The translator, therefore, produces a TL text through a *modification* or transformation of the original text.[3] This way of describing translation, which does not explicitly require a complete understanding of the SL text, is especially relevant when building machine translation systems.

Syntactical transfer systems (see figure 2, compare with figure 1) stem from an additional approximation: as many professional translators, they do not need to *understand* the sentences. Transfer systems go from SL parse tree and words directly into TL parse tree and words by applying parse tree transformations (*structural transfer*) and word substitutions (*lexical transfer*) without building an explicit representation of the interpretation of SL sentences. Lexical transfer and structural transfer are often designed in such a way that they are performed rather independently from each other, a fact that constitutes an additional approximation.

2.6 Additional approximations

Transformer systems (Arnold et al., 1994, 4.2), many of them commercial and available on the internet[4], may be seen as a more radical approximation which substitutes *full parsing* and tree

[3]Sager (1993, p. 116) defines translation as "a range of deliberate human abilities [...] which consist in text production in a target language, based, inter alia, on the modification of a text in a source language to make it appropriate for its intended new purpose".

[4]For example, SDL Transcend is available through `http://www.freetranslation.com` and Reverso is available as `http://www.reverso.net`.

SL sentence
↓
Analysis
↓
SL parse tree + SL words
↓
Transfer (syntactical and lexical)
↓
TL parse tree + TL words
↓
Generation
↓
TL sentence

Figure 2: Translating a sentence without explicitly building a representation of its interpretation.

transformation by chunk detection and chunk word reordering, where a chunk is a flat (unstructured) pattern of categories such as *article–adjective–noun*.[5]

2.7 Reliance on explicit knowledge

Translators have to be aware that, to design a MT system, their intuitive knowledge (Sprachgefühl) is useless because an elicitation of translation procedures (after a careful reflection) in the form of rules is required. In addition, only those rules that can be programmed in reasonable time and with reasonable effort into a MT system and only those which execute taking reasonable time and memory will be useful. These requirements are another source of approximations, sacrifices and compromises.

[5]For details on how to learn details of the particular rules of this kind of transformer architectures in the classroom, see Forcada (2000) or Mira i Gimènez and Forcada (1998).

2.8 Learning to view real MT as a crude approximation

Descending from the ideal model down to real systems may help translators to understand MT architectures in a theoretical setting, but, perhaps equally importantly, to become aware that efficient mechanization of a workable system implies very strong approximations —some of which apply also to human translation— and may help them to resituate their expectations about real MT: ideal MT is impossible and real MT cannot be programmed from intuitive knowledge: "this is why you cannot expect so much from real MT systems".

3 Refining "model zero" into a MT strategy

3.1 "Model zero"

Model zero reads the SL sentence word by word and produces a TL "sentence" by stringing constant, context-independent equivalents of SL words in exactly the same order (this should not be confused with vague concepts such as "literal translation"). For example if the SL text is $s_1 s_2 s_3 \ldots s_N$ where s_i $(1 \leq i \leq N)$ is a word, the TL text is $T(s_1)T(s_2)T(s_3)\ldots T(s_N)$ where $T(s_i)$ is the TL word chosen to be the fixed equivalent of s_i.

3.2 Obvious errors and intuitive refinements

Model zero produces a number of consistent errors, the most obvious being:

1. incorrect choice of equivalents for lexically ambiguous words (e.g., for homographs);

2. incorrect TL word order when SL word order and TL word order should be different;

3. agreement errors (the gender or number of a SL noun and its TL counterpart may be different and should change also for their modifiers);

4. wrong translations for idioms and other multi-word units.

After careful examination of the errors produced by a hypothetical model-zero system, and with some guidance, students readily propose intuitive partial solutions for the above:

1. including some kind of homograph resolution module, based on the lexical categories of adjacent words (this requires a morphological analysis module);

2. adding simple rules to reorder certain sequences of words according to their lexical categories;

3. adding agreement rules for certain sequences of words, again according to their lexical categories;

4. treating fixed-length multi-word expressions as single words (maybe managing their inflection through morphological analysis).

This intuitive model is actually not very far from the one used by real, commercial *transformer systems* as discussed in 2.6.

3.3 Comparing real MT to model zero

A laboratory assignment (Pérez-Ortiz and Forcada, 2001) may illustrate how real MT departs from model zero by studying the differences between the translation produced by a commercial MT system, first for each word in isolation and then inside a sentence, and hypothesizing what additional operations the MT system is doing besides substituting words: disambiguation, reorderings, agreement adjustments, multiword units, etc. Transformer architectures may then be studied as an improvement over word-for-word translation.

3.4 Viewing real MT as an improvement over word-for-word translation

Ascending from model zero up to real MT helps also shape the expectations of translators about MT. Real MT is not word-for-word, and is quite far from word-for-word MT: "to get to what one might call poor MT, a lot has to be done beside substituting words".

4 Conclusion

Placing real MT between an ideal, linguistically motivated model (as an inevitable approximation of it) and a radical word-for-word model called *model zero* (as a necessary refinement over it) may be a very powerful teaching tool to illustrate MT strategies in the classroom as well as to change the expectations of translators about real MT into a more realistic setting (see figure 3).

References

Allen, J. (1995). *Natural Language Understanding*. Benjamin/Cummings, Redwood City, CA, 2nd. edition.

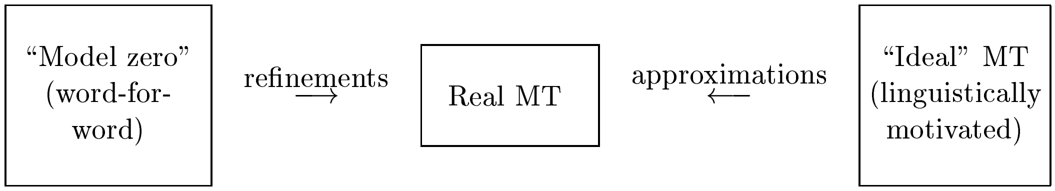

Figure 3: Real MT seen both as a refinement of a word-for-word model and as an approximation to an ideal, linguistically motivated, compositional model.

Arnold, D., Balkan, L., Meijer, S., Humphreys, R., and Sadler, L. (1994). *Machine translation: An introductory guide*. NCC Blackwell, Oxford. (available at `http://clwww.essex.ac.uk/~doug/MTbook/`).

Forcada, M. L. (2000). Learning machine translation strategies using commercial systems: discovering word-reordering rules. In Lewis, D. and Mitkov, R., editors, *MT 2000: Machine Translation and Multilingual Applications in the New Millenium*, pages 7–1–7–8, University of Exeter, UK. (available as `http://www.dlsi.ua.es/~mlf/docum/forcada00p.pdf`).

Mira i Gimènez, M. and Forcada, M. L. (1998). Understanding PC-based machine translation systems for evaluation, teaching and reverse engineering: the treatment of noun phrases in Power Translator. *Machine Translation Review (British Computer Society)*, 7:20–27. (online: `http://www.bcs.org.uk/siggroup/nalatran/mtreview/mtr-7/nalamt78.htm`).

Mori, T., Matsuo, M., and Nakagawa, H. (1999). Zero-subject resolution using linguistic constraints and defaults: The case of Japanese instruction manuals. *Machine Translation*, 14(3-4):231–245.

Nakaiwa, H. (1999). Automatic extraction of rules for anaphora resolution of Japanese zero pronouns in Japanese-English machine translation from aligned sentence pairs. *Machine Translation*, 14(3-4):247–279.

Pérez-Ortiz, J. A. and Forcada, M. L. (2001). Discovering machine translation strategies beyond word-for-word translation: a laboratory assignment. In Forcada, M., Pérez-Ortiz, J., and Lewis, D. R., editors, *MT Summit VIII, Proceedings of the Workshop on Teaching Machine Translation*. (available at `http://www.dlsi.ua.es/tmt/`).

Radford, A., Atkinson, M., Britain, D., Clahsen, H., and Spencer, A. (1999). *Linguistics: An introduction*. Cambridge Univ. Press, Cambridge.

Sager, J. C. (1993). *Language engineering and translation: consequences of automation*. Benjamins, Amsterdam.

Tellier, I. (2000). Semantic-driven emergence of syntax: the principle of compositionality upside-down. In *Proc. 3rd Conference on the The Evolution of Language*, pages 220–224, Paris.