The (Un)faithful Machine Translator

Ruth Jones
Dept. of French and Francophone Studies
University of California Los Angeles

Ann Irvine
Center for Language and Speech Processing
Johns Hopkins University

Abstract

Applying machine translation (MT) to literary texts involves the same domain shift challenges that arise for any sublanguage (e.g. medical or scientific). However, it also introduces additional challenges. One focus in the discussion of translation theory in the humanities has been on the human translator's role in staying faithful to an original text versus adapting it to make it more familiar to readers. In contrast to other domains, one objective in literary translation is to preserve the *experience* of reading a text when moving to the target language. We use existing MT systems to translate samples of French literature into English. We then use qualitative analysis grounded in translation theory and real example outputs in order to address what makes literary translation particularly hard and the potential role of the machine in it.

1 Introduction

The question of how to translate, especially when the source text is valued for its perceived literary merit, has been the focus of a discussion that is nearly as old as written text itself. A key debate is whether the translator should (1) adapt the source language text as it is translated into the target language to make it familiar and understandable to the reader, or (2) stay as faithful as possible to the original. Schleiermacher (2012) calls the former a *free* translation and the latter *faithful*. The former has also been referred to as *domesticating* the text, or *bringing the text to the reader*, in contrast to *foreignizing* the text, or *bringing the reader to the text* (Venuti, 2008; Berman, 1992).

Consider the French phrase *enculer les mouches*. Staying as faithful to the original French as possible, the first word, *enculer* translates as the infinitive for the French word for anal penetration, while the second is the more banal *flies*. Google translate gives *to fuck flies*. However, idiomatically, it is, despite the strongly sexual first term, a not uncommon way to say *to nitpick*. This translation makes the text more understandable, at the cost of staying faithful to the meanings of the individual words of the original text. Stylistic elements such as metaphor, alliteration, metonymy, and rhyme likewise require the translator to make interpretive choices beyond the literal meaning of the original, bringing the original text to the reader of the translation even at the expense of losing some of the literal meaning of the source.

Often multiple equally faithful translations of a word or phrase exist, and the translator must choose one based on context, either local or more broad. For example, the French *il neige* can be translated as *it snows* or *it is snowing*.[1] In English, *it is snowing* suggests the narrative present, while *it snows* suggests a habitual occurrence.

Like human translators, a statistical machine translation (SMT) system may produce translations that are relatively free or faithful and must constantly make translation choices in decoding. For SMT, choices are dependent on what is observed in training and language modeling data and their frequencies. When systems are trained on datasets that are similar to a test text, they are more likely to make reasonable translation choices. Additionally, if a model, either a priori or automatically, knows something about what the output should look like (e.g. poetry should rhyme or have rhythm), features could encourage free translations to take a certain form.

How much a translation sounds like an original text in its target language and how much it preserves elements of its source language, which make it sound foreign, is in part an ethical choice made by the human translator. Still, even experienced human translators have difficulty recogniz-

---

[1] There is no present progressive tense in French.

ing when they are being faithful and when their cultural experiences have influenced a translation. Current SMT models have no awareness of this and no ability to make specific choices to balance the two tendencies in the same output. Our work shines a light on SMT from the perspective of translation theory based on a qualitative analysis of two translated samples of French literature, one prose and one poetry. We compare SMT and human translations to address the following:

- What types of translation choices does the machine make, compared with humans?
- Is there evidence for the need to encourage a machine to translate more freely?
- Can SMT translate non-ethnocentrically?

2 Background

2.1 Translation Theory

Schleiermacher (2012) raises the issue of a translation's approximation of its source language vs. its fluency or resemblance to an original work in its target language, referring to translations "that are faithful or free." Berman (1992), alternatively, outlined the need for an ethics and an analytics of translation. For Berman, the translator has an imperative to avoid "freedom" where it brings with it a tendency to alter the foreign text by making it resemble a work of literature created in the target language through adjustments to the original on the levels of style, idiom, and content (both lexical and explicative). His is an argument for what Venuti (2008) calls "foreignization" in translation, preserving the distance between the language of the original text and the language of the translation by creating a translation that is perceptibly different from an original work in the target language. He opposes this to domestication, which instead privileges fluency and readability.

Venuti (2008) uses a similar critique to address the relative visibility or invisibility of the translator. For Venuti, part of the domestication of the translated text comes in the form of the invisibility of its translator in the finished (marketed) product. Compare, for instance, Venuti's example of the translator's invisibility in the 2002 Penguin translation of the Anna Karenina, advertised with praise for the "transparency" of its translation without naming the translators, to Seamus Heany's 2000 translation of Beowulf, which includes both original and translated texts side-by-side and features the poet/translator's name prominently on the cover. In the first case, the reader is asked to forget that she is not, in fact, reading Tolstoy in his own words, while, in the second, Heany's text is open to constant comparison with its original.

2.2 MT of Non-Standard Language

Prior work applying SMT to non-standard language focuses primarily on domain adaptation. In that task, an MT system trained on, for example, newswire, is used to translate text in a different domain, such as science. Much of this work has focused on up-weighting subsets of the training or language modeling data that are most similar to the new domain (Matsoukas et al., 2009; Foster et al., 2010; Ananthakrishnan et al., 2011; Niehues and Waibel, 2010; Foster and Kuhn, 2007; Tiedemann, 2010; Lavergne et al., 2011).

Other work has focused on literary texts (Reddy and Knight, 2011; Kao and Jurafsky, 2012; Roque, 2012). Most relevant is Greene et al. (2010), which presents a model for translating Italian poetry into English. That work focuses on preserving meaning as well as rhythm and is an interesting first attempt at integrating models of poetry ("how to say") and storyline ("what to say") generation. In many cases, it is hard to do both well at once; simultaneously maintaining the meaning and rhythm of a poem is challenging.

3 Experiments

3.1 Data and Setup

We analyze translations of two samples of French literature, one prose and one poem (Figures 1-2). The prose selection is a sample of the twentieth century novel L'Étranger by Albert Camus (Camus, 1955). We use the Camus and Ward (1989) English translation as a reference. The poetry selection is a sample of the twentieth century poem "Jardin" by Yves Bonnefoy (Bonnefoy, 1968), from the collection *Début et fin de la neige*, translated in Bonnefoy et al. (2012). We selected the passages because they use fairly simple language and have modern and well-known authors.

We translate the two literary selections using two SMT systems. First, we train a phrase-based MT model using the Hansard data.[2] The corpus contains over 8 million parallel lines of text and is one of the largest freely available parallel corpora for any language pair. It contains proceedings of the Canadian parliament. Recent work has

---

[2]http://www.parl.gc.ca

shown that newswire corpora, the other common bitext domain, is not very different from the parliamentary domain. Thus, a model trained on the Hansard data reflects the status of a typical modern SMT system trained on freely available data. We use the Moses SMT framework (Koehn et al., 2007), GIZA++ automatic word alignments (Och and Ney, 2003), and the batch version of MIRA for tuning (Cherry and Foster, 2012). For comparison, we also present and analyze translations by Google translate.[3]

In addition to our detailed manual analysis, we automatically evaluated outputs using case-insensitive BLEU and a single reference. The Moses system achieve a slightly higher BLEU score than Google (16.62 vs. 11.25) on the Bonnefoy selection and the opposite is true for the Camus selection (26.03 vs. 30.05). However, because the selections are small, we don't interpret these results as particularly meaningful.

3.2 Analysis

Figures 1 and 2 show the outputs. Focusing on the differences between the machine and human translations with respect to the originals reveals places where the human translators' choices diverged from the translations considered probable by the two SMT models. Close reading of the source text and human translation suggests possible reasons for the translator's choices. The probabilities that the SMT model assigns to the human translations relative to those assigned to the observed MT output highlights the need for probabilistic translation models that are specific to the domain of the particular texts or for literary translation more generally. While differences occurred based on a variety of factors, for the sake of brevity, we only consider lexical variation and the question of time as an aspect of translation. We take examples from Camus' prose and Bonnefoy's poem while keeping in mind the possibility of a definable difference in domain between the two.

Ward's translation puts the plain language of Camus' text into a clear and conversational English, a stylistic choice for fluency and domestication of the French. The focus in this passage is on the speaker, actively calming himself after the departure of the warden, and the "re" of "retrouvé" appears as "again." The machine translations, looking at the words in the absence

---
[3] translate.google.com

L'étranger - Albert Camus

lui parti, j' ai retrouvé le calme.
j' étais épuisé et je me suis jeté sur ma couchette.
je crois que j' ai dormi parce que je me suis réveillé avec des étoiles sur le visage.
des bruits de campagne montaient jusqu' à moi.
des odeurs de nuit , de terre et de sel rafraîchissaient mes tempes.
la merveilleuse paix de cet été endormi entrait en moi comme une marée.
a ce moment, et à la limite de la nuit, des sirènes ont hurlé.
elles annonçaient des départs pour un monde qui maintenant m' était à jamais indifférent.
pour la première fois depuis bien longtemps j' ai pensé à maman.

Hansard SMT

his party, i found the calm.
i was exhausted and i had on my sleeper.
i think i slept because i woke up with stars on your face.
the campaign noise up to me.
the smell of night of land and salt rafraîchissaient my tempes.
the great peace of this summer asleep came into me as a tide.
has this time, and the limit of the night of sirens screamed and yelled.
they announced departures for a world that now was never indifferent.
for the first time in a long time i thought about mom.

Google SMT

he was gone, i found calm.
i was exhausted and i threw myself on my bunk.
i think i slept because i woke up with stars on her face.
noises campaign amounted to me.
the smell of night, earth and salt refreshed my temples.
heavenly peace this summer sleeping entered me like a tide.
at that time, and the limit of the night, sirens screamed .
they announced departures for a world that now was never indifferent to me.
for the first time in ages i thought mom.

The Stranger, translation by Matthew Ward

with him gone , i was able to calm down again.
i was exhausted and threw myself on my bunk.
i must have fallen asleep, because i woke up with the stars in my face.
sounds of the countryside were drifting in.
smells of night, earth, and salt air were cooling my temples.
the wondrous peace of that sleeping summer flowed through me like a tide.
then, in the dark hour before dawn, sirens blasted.
they were announcing departures for a world that now and forever meant nothing to me.
for the first time in a long time i thought about maman.

Figure 1: The Stranger by Albert Camus

of Camus' protagonist, give "found," eliminating the "re." Ward translates "se calmer" exactly, "to calm (down)." In contrast, the machine versions give "found (the) calm." It is not the passive aspect of Camus' phrase that is problematic ("finding calm" as opposed to "calming down"); rather, it is the return implied by the "re" that gives pause. Ward's translation gives a plainer, more informal style than the translations offered by the SMT systems, choosing to preserve the repetition of "re" (in "retrouvé") with "again" rather than the core meaning of "found" in "trouvé."

Later in the passage (line 3), the phrase "je

Début et fin de la neige, Yves Bonnefoy, « Le jardin »
il neige. sous les flocons la porte ouvre enfin au jardin de plus que le monde. j' avance. mais se prend mon écharpe à du fer rouillé, et se déchire en moi l' étoffe du songe. il neige. sous les flocons la porte ouvre enfin au jardin de plus que le monde. j' avance. mais se prend mon écharpe à du fer rouillé, et se déchire en moi l' étoffe du songe.

Hansard SMT
it snows. under the snowflakes the door opens finally au jardin more than the world. but is my point. my scarf to iron rusty tears, i think in character. it snows. under the cornflakes and opens the door au jardin de more than the world. my point. but does my scarf to iron rusty, that tears character in me thinking.

Google SMT
it snows. flakes under the door finally opens to the garden over the world. i advance. but takes my scarf with iron rusty, and tears in me the stuff of dreams. it snows. finally, in the snow the door opens to the garden over the world. i advance. but take my scarf of rusty iron, and tears in me the stuff of dreams.

Beginning and End of the Snow, Emily Grolsholz, "The Garden"
it's snowing. beneath the snowflakes the gate opens at last on the garden of more than the world. i enter. but my scarf catches on rusty iron, and it tears apart in me the fabric of the dream. it's snowing. beneath the snowflakes the gate opens at last on the garden of more than the world. i enter. but my scarf catches on rusty iron, and it tears apart in me the fabric of the dream.

Figure 2: The Garden by Yves Bonnefoy

me suis réveillé avec des étoiles sur le visage" is translated as "I woke up with the stars in my face" in Ward's translation, whereas the Hansard and Google translations drop the indefinite article and assume a second person in the scene, giving "i woke up with stars on {your, her} face." Later, the phrase "des bruits de campagne" (line 4) also provides a source of linguistic confusion.

It is "sounds of the countryside" in Ward, but "the campaign noise" and "noises campaign" in Hansard and Google, respectively. Ward's translations make two distinct choices for the indefinite article "des," converting it to a definite article (the) in the first instance while dropping it in the second. Both examples again show Ward working the text into plain-spoken English prose by choosing the specific "the stars" over the general "stars" for "des étoiles" and the more conventional construction sounds of the coutryside over countryside sounds, which would preserve the unfamiliar (as shown by the difficulty of both MT systems in translating this phrase) construction of "des bruits de campagne." The discrepancies between the human and MT versions of Camus' text suggest that the MT systems might, at the least, be able to identify the difficulties of translating certain stylistic elements of the French.

The translations of Bonnefoy's poem reveal slightly different concerns. The translations of "étoffe" exemplify a lexical choice problem. Grolsholz's choice of "fabric" has a lower translation probability in the SMT models than "stuff" (Google translation). Both meanings are possible, but while "stuff" is more common, the source text suggests an association between "écharpe" (scarf) and "étoffe" (stuff/fabric) that comes to the fore in Grolsholz's translation. Taken with similar choices ("gate" for "door", also "snowflakes" for "flocons," earlier in the poem), Grolsholz's translation reveals a preference for specificity over probability that goes beyond rhythmic consistency to effect the translated poem's recreation of the images present in the original.

Temporality also appears as a difference between Grolsholz's and the machine translations. Specifically, Grolscholz translates "il neige" (line 1) as "it is snowing." Neither SMT model selected the present progressive. Their translation, "it snows" has a distinctly high probability in the Hansard model, as the parliamentary proceedings deal most often with general conditions when discussing weather (i.e. "it snows in the prairies"). While this is an adequate translation of the French phrase, Grolsholz's choice of the progressive anchors the poem in a narrative present that is absent in the general expression "it snows." This moment is key to understanding the poem in the context of the larger collection, as it gives the poet a defined position in time that anchors the poem's imagery.

The fact that neither MT system made this choice suggests a difference between literary and nonliterary texts in terms of how each treats time and the experience of duration. Temporality functions in subtly different ways in French and English. It is important to narrative and literary text and is particularly difficult for the MT system.

4 Discussion

Defining the type and degree of domestication that a literary translation should take is difficult to express, even to a human. We can say that Ward's translation, with its conversational style and choice of sense and style over language play, is more domestic than Grolsholz's, which tries to reflect the syntax of the original. Indeed, if we look back to Venuti's complaint about the translation of Anna Karenina, Grolsholz is certainly the more visible of the two translators, each of her translations being accompanied by its original on the facing page. From a technical standpoint, we may want a translation to take into consideration the narrative of a text in order to describe events in the narrative present (e.g. choosing "it is snowing" over "it snows"). However, defining the scope of the relevant narrative context is difficult and may vary substantially from text to text.

From the ethical perspective of the foreign/domestic debate, deciding how much the narrative context needs to be explicated or altered to be understandable in the translation is dependent on variables including the translator's stance on this issue, the author's wishes (if the author is living) and the publisher's requirements. Even once they have been determined, specifying such preferences precisely enough for a computational model to follow is even harder. For example, we could model a general preference for specific translations of nouns over more probable translations (e.g. 'snowflakes' instead of 'flakes'), but translation rules are typically very noisy and an SMT system would likely be tempted by garbage translation rules (e.g. in the Hansard system, 'flocons' translates as 'cornflakes' with higher probability than 'snow', 'flakes', or 'snowflakes'). In short, part of the human translator's job is knowing when to make exceptions to convention for the sake of the reader's experience of the translated text, and the question of the exception is difficult for the machine to account for.

Even if the type and degree to which a text should be domesticated could be accurately modeled, some types of free/fluent/flexible translations will be easier for a machine to produce than others. For example, idioms may be easy to integrate; if they are observed in training data, then a machine can easily produce them. This, however, requires in-domain training data, and domain is somewhat of a moving target in literature due to extremely high variability. In contrast to the ease of memorizing static idioms, computationally choosing correct, relevant, and appropriately specific translations of individual nouns (e.g. 'porte' as 'gate' instead of 'door') is difficult.

We end our discussion on a note about *visibility*. Introducing an SMT system into debates surrounding literary translation by human translators would seem to cause the translator to disappear entirely. Indeed, according to Cronin (2012), "machine translation services would appear to render invisible the labour of translation..." However, for Venuti, visibility is crucial to the ethics of balancing domestication and foreignization to create non-ethnocentric translations in that it reminds the reader to be attentive to the translation and to the translator as creative labourer. As a level of domestication is to be expected in fluent translations, Venuti's argument for visibility is also an argument for a disruption to the reader's experience that reinserts the distance of the foreignizing translation in a different way, suggesting that fluency, which hides the act of translation, might be ethical under conditions of visibility. Difficulties encountered by an SMT system can constitute a kind of visibility, because they expose problems in the translation, which often come in the form of disfluencies. However, these systems cannot consider translation in terms of domestication and foreignization; the SMT objective is to use patterns observed in training data example translations to produce something that has the same meaning as the source text and looks like the target language. There is a constant tradeoff between fluency and faithfulness. Although SMT can deal with fluency, it cannot handle ideas of domestic and foreign. Therefore, if we accept that domesticating and foreignization is key to distinguishing visibility, then the relationship between visibility and invisibility for the human translator and the machine translator must be different. And this divergence, in turn, means that current approaches to SMT could not ensure non-ethnocentric translations.

References

Sankaranarayanan Ananthakrishnan, Rohit Prasad, and Prem Natarajan. 2011. On-line language model biasing for statistical machine translation. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*.

Antoine Berman. 1992. *The Experience of the Foreign: Culture and Translation in Romantic Germany*. State University of New York Press, New York. Trans. by S. Heyvaert.

Y. Bonnefoy, E. Grosholz, and F. Ostovani. 2012. *Beginning and End of the Snow / Debut Et Fin de la Neige: Followed by Where the Arrow Falls / Suivi de La Ou Retombe la Fleche*. Rowman & Littlefield Publishers, Incorporated.

Yves Bonnefoy. 1968. *Selected poems*. Cape editions. Cape.

A. Camus and M. Ward. 1989. *The Stranger*. Everyman's library. Knopf Doubleday Publishing Group.

A. Camus. 1955. *L'étranger*. Appleton-Century-Crofts.

Colin Cherry and George Foster. 2012. Batch tuning strategies for statistical machine translation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Michael Cronin. 2012. *The Translation Age: Translation, Technology, and the new Instrumentalism*. The Translation Studies Reader, Third Editions. Routledge, New York.

George Foster and Roland Kuhn. 2007. Mixture-model adaptation for smt. In *Proceedings of the Workshop on Statistical Machine Translation (WMT)*.

George Foster, Cyril Goutte, and Roland Kuhn. 2010. Discriminative instance weighting for domain adaptation in statistical machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Erica Greene, Tugba Bodrumlu, and Kevin Knight. 2010. Automatic analysis of rhythmic poetry with applications to generation and translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Justine Kao and Dan Jurafsky. 2012. A computational analysis of style, affect, and imagery in contemporary poetry. In *NAACL Workshop on Computational Linguistics for Literature*.

Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*.

T. Lavergne, A. Allauzen, H. Le, and F. Yvon. 2011. LIMSI's experiments in domain adaptation for IWSLT11. In *Proceedings of the International Workshop on Spoken Language Translation (IWSLT)*.

Spyros Matsoukas, Antti-Veikko I. Rosti, and Bing Zhang. 2009. Discriminative corpus weight estimation for machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Jan Niehues and Alex Waibel. 2010. Domain adaptation in statistical machine translation using factored translation models. In *Proceedings of the European Association for Machine Translation (EAMT)*.

Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Comput. Linguist.*, 29(1):19–51, March.

Sravana Reddy and Kevin Knight. 2011. Unsupervised discovery of rhyme schemes. In *Proceedings of the Conference of the Association for Computational Linguistics (ACL)*.

Antonio Roque. 2012. Towards a computational approach to literary text analysis. In *NAACL Workshop on Computational Linguistics for Literature*.

Fredreich Schleiermacher. 2012. *On different methods of translating*. The Translation Studies Reader, Third Editions. Routledge, New York. Trans. by Susan Bernofsky.

Jörg Tiedemann. 2010. To cache or not to cache? experiments with adaptive models in statistical machine translation. In *Proceedings of the ACL Workshop on Statistical Machine Translation and Metrics (MATR)*.

Lawrence Venuti. 2008. *The Translator's Invisibility: A History of Translation*. Routledge, New York.