

Title	Topsy-Turvydom: Does AI Have a Place in a Japanese-English Translation Course for University Students?
Sub Title	トプシイタービイダム : 英訳講座の受講生のAI仲間いり？
Author	Armour, Andrew
Publisher	慶應義塾大学日吉紀要刊行委員会
Publication year	2019
Jtitle	慶應義塾大学日吉紀要. 英語英米文学 (The Hiyoshi review of English studies). No.71 (2019. 3) ,p.1- 26
JaLC DOI	
Abstract	
Notes	
Genre	Departmental Bulletin Paper
URL	https://koara.lib.keio.ac.jp/xoonips/modules/xoonips/detail.php?koara_id=AN10030060-20190331-0001

慶應義塾大学学術情報リポジトリ(KOARA)に掲載されているコンテンツの著作権は、それぞれの著作者、学会または出版社/発行者に帰属し、その権利は著作権法によって保護されています。引用にあたっては、著作権法を遵守してご利用ください。

The copyrights of content available on the Keio Associated Repository of Academic resources (KOARA) belong to the respective authors, academic societies, or publishers/issuers, and these rights are protected by the Japanese Copyright Act. When quoting the content, please follow the Japanese copyright act.

Topsy-Turvydom: Does AI Have a Place in a Japanese-English Translation Course for University Students?

Andrew Armour

Abstract: The focus of this paper is a course in Japanese-English translation taught for the Research Center for Foreign Language Education at Keio University by the author for the past fourteen years. The practices developed over this period are explained along with the difficulties faced by the students, many of whom are Japanese. Since translation is best done into one’s native language, it can be said these students are approaching it “topsy-turvy”. However, this challenge affords them a deeper understanding of the ways in which the Japanese express themselves and how substantially these differ from English. Unlike French-English translation, simple replacement of Japanese words and phrases with their English equivalents in the dictionary often is not productive, something which is borne out by phrase-based translation software. Machine translation has been used in the course since 2006; the extensive errors that resulted proved instructive as well as entertaining. However, the adoption of neural machine translation (NMT) technology for Google Translate in late 2016 changed the playing field.¹⁾ AI can now serve as worthy competition for students in a JE translation course, allowing them to gauge their own weaknesses and strengths while honing their skills.

Introduction

The term “topsy-turvy” – in the sense of an inversion of the natural or proper order – has long been applied to Japan, its culture and language. It

was, for instance, the title of Mike Leigh's 1999 film about Gilbert and Sullivan and their 1885 comic opera entitled *The Mikado*. In fact, the association long predates the Victorian period²⁾ and continues to the present day, as demonstrated by an observation on traditional Japanese woodworking techniques in a 2017 documentary:³⁾ in the West the carpenter *pushes* both saw and plane, while in Japan these tools are *pulled*. Whatever the merits of such *nihonjinron*, the reason I have chosen to use the term “topsy-turvy” is to underscore how AI is turning the tables and transforming the teaching of translation.

In the early days of machine translation (MT), a field which developed from the 1950s, it was linguists and translators who helped to train the computer systems being developed for the automated translation of texts in the military, commercial and even religious fields. Progress was also dependent on the hardware, which rapidly grew more powerful and affordable. In the 1980s, with the advent of the personal computer and word-processing software, human translators were freed from the tyranny of the typewriter and gradually came to benefit from Japanese digital dictionaries.⁴⁾ Meanwhile, Japanese companies developed and marketed expensive translation software for business use; their focus was the translation of English to Japanese, rather than the reverse. It may be questioned whether investment in such programs was ever worthwhile. Before long, however, fresh competition appeared in the form of free online translation services for short texts, such as that offered in 1996 by SYSTRAN (AltaVista Babelfish). The quality was perhaps not much inferior to that of the commercial software. Nevertheless, in either case MT was still only useful for “gisting” – that is, providing the sense of a source text – and even for that was not reliable.

It was generally acknowledged that, until a quantum leap was achieved

in MT, the only practical use of computers for translation was to operate them in tandem with human editors and translators – that is, computer-aided translation (CAT). This usually means that a translator post-edits the text output by a computer, although pre-editing may also be performed to remove ambiguities and thus make the computer’s task easier. Dominant in the CAT field are translation memory (TM) tools, such as SDL Trados, that rely on databases of matched segments of text in Japanese and English. Thus, if the TM software is presented with a Japanese text including the phrase 「チョコレートに目がない」 which is already paired with “I love chocolate” in the database, it will insert the English in place of the Japanese.⁵⁾ All recognized segments are replaced in this way, leaving the human translator to (a) check the output, correcting any errors, and (b) translate segments which could not be matched. The human input is then fed back into the database in a process of continuous refinement. I have no experience of TM software, although I believe it could perhaps play a role in the classroom. What I have made use of, however, for teaching is general-purpose MT software, free on-line services such as Babelfish and Reverso, and more recently Google Translate.

Screening

Since 2005 I have run a Japanese-English translation course on Keio University’s Mita campus that is attended predominantly by Japanese students from various faculties. Although there are rarely more applicants than places for this course, at the start of the academic year they are required to take a test designed to gauge their general familiarity with the target language, English. Typically, there are three parts to this: (1) a paragraph in which all vowels have been deleted, (2) a paragraph in which all punctuation and capitalization has been removed, and (3) a paragraph in which the order

of the sentences has been scrambled. The first two are unfamiliar to most Japanese university students and serve to determine their degree of fluency. The third part focuses on logical organization, which is important when translating into English. During the course, for example, the students are frequently encouraged to reposition sentences and even whole paragraphs to produce a more coherent flow.

Only students who pass the screening may take the course. In the case of foreign students with native or near-native fluency in English, a separate test is used to confirm their ability to comprehend spoken/written Japanese.

Orientation

In the week following screening, the students are provided with some guidelines regarding how they should approach JE translation.

Translation involves decoding the meaning and nuances of the source text and encoding this information into the target language. Needless to say, the first process requires an intimate knowledge not only of the language but also of the culture. Here the Japanese student may feel confident. There are pitfalls, however, such as false friends: Japanese students will invariably assume that an English word that has become established in Japanese and written in *katakana* need only be converted back into the alphabet. This is so often *not* the case that I advise students to check *every* time.

The second process, re-encoding the meaning of the original into English, is the real challenge for most students, despite being semi-fluent in English. Nevertheless, there are ways in which they can approach the task with more confidence.

Firstly, I impress on them that translation is facilitated by curiosity and imagination. Curiosity is important if they are to accumulate general knowledge so as to be able to place information in context even before they

start any research. This means keeping up with the news in a variety of fields – politics, economy, technology, etc. Imagination is often required in order to fill in any lacunae – gaps in the source text, owing to insufficient information or reasoning – and to clarify the ambiguities so often found in a source text. It helps if the students can put themselves in the author’s shoes, to adopt the perspective of the writer, who is rarely available to answer questions like “What did you mean by this?”

Secondly, when tackling an assignment, the student should read around the topic to become an “insider”. This rarely means expanding their research beyond *Wikipedia*, but the principle is to become sufficiently familiar with the topic to be able to add, subtract and rearrange the content with some degree of confidence. The most frequently repeated advice I give my students is “Forget the Japanese!”,⁶⁾ and insider knowledge is what helps to make this possible.

Some students are more adept at background research than others, but all should share the same objectives, so I set out what I believe should be the students’ goals:

- To make a translation look like an original
- To add what is necessary to ensure clarity
- To subtract everything that is not necessary
- To rearrange the text so the order is (chrono)logical
- To create, if possible, something better than the original

The first and last of these would appear to be asking too much of a student whose first language is not English. However, it is important to set the bar high if they are to make appreciable progress.

The students are also presented with a recommended procedure for

tackling a translation assignment:

1. Read at least twice and thoroughly grasp the content.
2. Create a draft translation, highlighting any possible problems.
3. Check that all significant information has been transferred – accurately.
4. Hide the Japanese and polish your English translation.
5. Reread and add “spice” (preferably the following day).

The first four steps may seem obvious, but one of the main hurdles faced by the students is that they try to translate *while looking at the Japanese*. This should be avoided as much as possible, even for step 2, since it leads to the same sort of errors as produced by phrase-based MT software.

The final step is adding “spice”: a snippet of information, not found in the original but which adds extra clarity, or some terminology specific to the field. This can make all the difference: if well executed, the translation will look more like an original document. This is within the capabilities of a student whose first language is not English since it depends on research.

Additionally, I provide instruction on crafting a good title. The source text for an assignment is typically an extract, lacking a title. Translation necessitates close reading, and thus the students should be able to create a title based on their comprehension of not only the content but also the mood of the Japanese text.

Tools

In addition to using printed and electronic dictionaries (Eijiro and Weblio are recommended for translation use), students are urged to refer to a thesaurus. For Japanese students it can be difficult to know which word in a

thesaurus is preferred; in such cases, Google Ngram Viewer (books.google.com/ngrams) will show that, for example, “eruption” is much more common than “eructation”. A good alternative to a thesaurus is a dictionary viewer – that is, desktop software which enables one to look up a word in multiple dictionaries simultaneously. I demonstrate how to use DDwin with 20 EPWING-format dictionaries such as EDICT.⁷⁾ Sadly, DDwin is no longer maintained, and with the progress of machine translation it is unlikely that there will be further developments in this direction.⁸⁾

Advice

As already noted, the advice I repeat most is to “forget the Japanese.” When is this appropriate? Two examples should suffice: if the source text contains a noun (such as 見える化), the inexperienced Japanese student will typically look for an equivalent English noun, even though that may not be the best solution; and if the source text employs an adjective (such as 詳しく), they will similarly struggle to find the right English adjective, although a verb (e.g. “You know a lot about drones!”) may well be the best choice. Like a *précis*, a translation is best composed when not looking at the original text. The translator then draws on his or her own internal resources to render the sense of the original in their own (English) words.

The importance of both adding and subtracting when translating Japanese is repeatedly stressed throughout the course. It is not uncommon to have to supply missing background information if the writer assumed his readers were all Japanese and familiar with, say, Saigō Takamori or Sakurajima. More importantly, the translator frequently has to “join the dots”, especially when the content is not intended to be allusive. This is a practice that Japanese students take time to accept. Another hurdle for them is appreciating how paragraphs function in English; this knowledge – along

with the cognitive skills needed to rearrange text to create well-formed, linked paragraphs – is essential for producing academic translations.

Another piece of advice I give to my students is, if stuck, to adopt a sort of “topsy-turvy” approach – that is, to use the passive voice instead of the active, or a double negative to replace a positive statement – which is not something that comes naturally to them, even if they are capable of producing natural, colloquial English in conversation. Encouraging flexibility in their approach to a text is one goal of the course; for example, I introduce word puzzles such as 「春夏秋冬二升五合」, which has the added benefit of raising the topic of traditional weights and measures.⁹⁾

Materials

My intention has always been to select source materials that are widely varied in content and style.

Chronologically speaking, the oldest material is *Genji monogatari*; however, this is used not as an assignment but rather as a demonstration of the process by which non-Japanese translators, beginning with Arthur Waley (1889–1966), went about the process of turning a manuscript,¹⁰⁾ written with *hentaigana* and very few *kanji*, into flowing English prose. While Murasaki Shikibu’s tale was not the first work of Japanese literature to be translated into English, it was the first to become world famous. Indeed, Waley’s translation, though incomplete, was itself translated into other languages. Other early translators of Japanese literature, many of them British, are also introduced – together with their idiosyncrasies. One notable example is Basil Hall Chamberlain (1850–1935), who translated the *Kojiki* in 1882. His Victorian sensibilities led him to switch, without warning, from English to Latin whenever he felt the content was distasteful or otherwise unsuitable for women and children (presumably unable to read Latin):

Having descended from Heaven onto this island, they saw to the erection of an heavenly august pillar, they saw to the erection of a hall of eight fathoms. Tunc quaesivit [Augustus Mas-Qui-Invitat] a minore sorore Augusta Femina-Qui-Invitat: “Tuum corpus quo in modo factum est?”¹¹⁾

This provides a good lesson in how cultural attitudes have changed regarding the appropriate stance and responsibilities of the JE translator.

As regards style, Waley’s translation is compared with that of Edward Seidensticker (1921–2007). Waley was a poet and his prose evokes the classical elegance of the Heian period, while Seidensticker, whose translations are mostly of 20th-century authors, wrote in a distinctly modern style, sometimes stripping away adjectives where Waley may have even added some. Unsurprisingly Waley’s translation is longer, something that Seidensticker was keen to point out. Of course, “lean” translations do have their place, but Japanese students need little encouragement in this regard.

Works of Japanese literature comprise only a small fraction of the coursework, but the students are frequently confronted with source texts dealing with uniquely Japanese content, such as *kamaitachi* (sickle weasel), *matagi* (mountain men) and *netsuke*. It should be noted that all three of these examples have their own English-language pages in *Wikipedia*; this means the students can read around the topic and discover related vocabulary and useful phrases. Also, such assignments typically require explanatory footnotes; acquiring the skill to compose these is yet another objective of the course.

Other assignments cover new technology and products, political developments, diplomacy, film reviews, experimental psychology, medical treatments, social issues, and so on. They are taken from printed or online

publications, newspapers and magazines, blogs and interviews.

In addition, the students are given opportunities to engage in “scanlation” – adding English to a Japanese manga from which the Japanese text has been scrubbed. To lighten the burden, I scan the manga and erase the Japanese text using Photoshop. The students are then provided with a PowerPoint file that includes both the original scan and the textless version; they only need to create textboxes to hold their English translation.¹²⁾ This assignment is, therefore, not especially demanding – although it may be challenging to pick short English words that will fit into the slim dialog bubbles contoured for vertical text – but it does provide a hands-on introduction to a practice that has been key to the dissemination of Japanese pop-culture overseas.

Going hand-in-hand with scanlation is fansubbing – that is, the addition of subtitles to video material such as Japanese animé by fans (amateurs). Despite the obvious copyright problems, this too has played a key role in the spread of modern Japanese culture,¹³⁾ made possible by the availability of free desktop software such as Aegisub.¹⁴⁾ The students are introduced to this subtitling software, which they can download and install on their own computers for home assignments. Although Aegisub has not been updated for four years, it is remarkably well designed: students can work on the subtitle text (with a spell checker and thesaurus) and adjust timing while viewing the video source and an audio spectrum display. They can also apply different styles for color differentiation of speakers, etc. The popularity of this software has meant that there is no shortage of manuals and tutorials,¹⁵⁾ and for this reason it is only briefly introduced in class.

Once the students have experimented with Aegisub, they are set an assignment – typically a video interview, from YouTube, of about 4 minutes in length. The simplest subtitle files (denoted by the .srt extension) may

contain only timing and text data; for this reason, the students are required to use the default file format for Aegisub (denoted by the .ass extension), featuring Advanced Substation Alpha syntax, which can include font and positioning information. The latter feature is convenient as broadcast Japanese video materials will often contain Japanese subtitles, perhaps not for the hard of hearing but rather for emphasis and audience stimulation.¹⁶⁾ The student can use Aegisub to place the English subtitles above or to one side of any Japanese text.

Having polished their translation and timings, the students deliver the resulting .ass file by email for screening¹⁷⁾ and discussion in class. The body of a typical subtitle file, when viewed in a text editor, looks like this:

```
Dialogue: 0,0:00:04.30,0:00:09.92,plain,,0,0,0,,Kasumi
        Arimura\N{\il}Little Maestro{\i0}
Dialogue: 0,0:00:11.82,0:00:30.00,interviewer,,0,0,0,,How
        did you prepare for the role?
Dialogue: 0,0:00:12.35,0:00:23.30,actress,,0,0,0,,I tried to
        stick to basic principles.
```

This sample includes three different styles (plain, interviewer, actress), which are defined in the file header; colors help to identify who is talking (especially useful if an interlocutor is off-screen). The start and end of a single subtitle can be adjusted by 1/100th of a second. While this precision is welcome, students are urged not to spend too much time refining the timings. However, the creation of colloquial English dialog with a natural rhythm¹⁸⁾ is the main aim of this soft-subbing project.¹⁹⁾ It is also designed to give the students confidence in a translation process that they may think lies outside their comfort zone.

Feedback

Scanlation and fansubbing are the exception. Most assignments involve translating several paragraphs of Japanese text into English. These are delivered via email and compiled into a single Word file, with the following components:

1. Source text (Japanese)
2. Model translation (mine)
3. Machine translation 1 (Japanese to English)
4. Machine translation 2 (English back into Japanese)
5. Student translations, in order of receipt, anonymized

The student translations are then marked using the highlighting function of Word: individual words or phrases that are deemed good (above average) are highlighted in green, while problematic words and phrases, grammatical errors, missing capitalization and punctuation are highlighted in yellow. This provides instant visual feedback for the students, who, it is hoped, will see the green/yellow ratio gradually increase over the academic year. In addition, symbols are used to indicate that something is missing (⊙). A series of bullets (⊙⊙⊙) indicates some factual information has been omitted. In the class following an assignment, the students take turns to translate each sentence afresh, while I provide feedback. The model translation is then screened and explained, followed by the machine translations.

The complete Word file, with highlighting, is then distributed to the students as a PDF file so they can review their translations. They are also encouraged to learn from the mistakes of others, including those in the

machine translations.

Machine translation

From the academic year starting in April 2006, I have included machine translations along with my own (model) translation when returning an assignment to the students. For each assignment, I would pick either commercial software such as ATLAS 翻訳パーソナル2005 or 本格翻訳, or sometimes an on-line service such as Reverso (www.reverso.net) to perform (1) Japanese-English translation, and (2) English-Japanese translation. Since 2017 I have used Google Translate almost exclusively. Because the same software is used to perform both steps, the final Japanese output should bear a strong resemblance to the original Japanese input. Generally, this has not been the case, giving rise to some amusement as well as some intriguing and instructional errors.²⁰⁾

How then has machine translation measured up? The following text is part of an assignment set in 2009:

一人の男と彼の犬

イギリスにはさまざま興味深いスポーツがテレビで観戦できる。面白さという点においてはこれに勝るものはない。これは、一人の男がまず登場する。この男は羊飼いである。であるから、汚れた労働着を着、ゴム長を履いて、手に長い杖を持っている。顔などもすすけたような日焼け顔というか、れっきとした田舎面をしている。彼の脇には、利発そうな牧羊犬が控えている。

This was translated by Reverso thus:

A dog of one man and boyfriend

The U.K. can watch a game of the sports that various interest is deep on TV. There isn't a thing superior to this in a point to be Omojiro. At first one man appears in this. This man is a shepherdess. Because I appear, and there is it, I wear dirty workwear and wear rubber boots and have a stick having a long in a hand. I do the respected country surface whether a face mystery says a sunburn face stained with soot. The sheep dog which seems to be clever waits on his side.

If one were to estimate a BLEU score²¹⁾ for the above, it would be close to zero; the “translation” is virtually incomprehensible. My model translation reads thus:

One Man and His Dog

The British can enjoy all sorts of fascinating sports on television, but there is none more interesting than this. It starts with a man – a shepherd. As you would expect, he has on grubby working clothes and wellington boots, and holds a long stick in one hand. His weathered visage looks almost sooty from long exposure to the sun. He has the unmistakable look of the country about him. At his feet waits an alert-looking sheepdog.

Not surprisingly, the students' translations were much closer to the model than to the MT output. Incidentally, clues that a machine was behind the first translation are (a) the insertion of the first person where the subject was not stated; (b) the appearance of Japanese words, romanized, where the machine could not find a match; and (c) the general lack of common sense. These telltale signs have, however, almost vanished with the adoption of neural machine translation (NMT).

In addition to entertainment, the reason for including these machines translations had originally been twofold. First, I hoped to discourage the use of machine translation. Students now have access to a variety of on-line translation services, and for some there must be a temptation to use these, if only to provide a rough basis for a more polished translation. Furnishing a machine translation should, it was assumed, serve to dissuade students from the outset.

Second, machine translations have, until recently, been of such poor quality that I hoped they would demonstrate that even Japanese students lacking confidence in their English could fairly easily outdo computers. MT would thus also serve as an encouragement, I believed. This situation has, however, changed since November 2016 when Google upgraded its translation software.²²⁾ For assignment feedback I now use Google Translate²³⁾ as it appears to have few contenders, at least among the free on-line offerings. Also, whatever the final target language, it first renders the source into English, making it especially suited for my translation course.²⁴⁾ It will be remembered that the commercial software originally used for the course was designed principally to do the opposite – translation from English into Japanese.

Competition

The key to the success of the upgraded Google Translate is a neural machine translation engine known as Google Neural Machine Translation (GNMT). What this replaces is the phrase-based approach in which phrases or individual words are translated and reassembled into sentences, while referring to various grammatical rules:

Machine translation services such as Google Translate have mostly used

a “phrase-based” approach of breaking down sentences into words and phrases to be independently translated. But several years ago, Google began experimenting with a deep-learning technique, called neural machine translation, that can translate entire sentences without breaking them down into smaller components. That approach eventually reduced the number of Google Translate errors by at least 60 percent on many language pairs in comparison with the older, phrase-based approach.²⁵⁾

My experience with MT over the many years I have been teaching translation is that the improvement is indeed “at least 60 percent.”

The source text is read by an encoder – a network of 8 layers, with 1,024 nodes (“artificial neurons”) per layer. The Japanese is converted into a list of vectors. The principle here is cumulative: each vector represents the meanings of all the words read so far in the sentence, so that the last is the longest. This approach has been likened to a human translator returning to the source text repeatedly for confirmation. I encourage the students to do this when creating their first draft, referring back to the source as many times as is required to confirm that all information and nuance is reflected in the draft translation, however rough it may be. But when the confirmation process is complete, they should “forget the Japanese” and focus entirely on the draft, polishing the English to convey the sense of the original, freely moving phrases, sentences and sometimes whole paragraphs if the end result is more lucid.

Traditional machine translation will treat words and phrases separately, so that context is lost. A “man” becomes a “shepherdess” before the end of the sentence. And even though the start of the sentence is about the elderly, that information has already been lost by the time the word ホーム is encountered, perhaps giving rise to the translation “platform” instead of the

correct “home”. In other words, machine translation has suffered from very poor short-term memory, giving human translators an advantage – until the application of deep learning.

With GNMT, another 8-layer network, the decoder, is responsible for generating the English translation. Further sophistication is supplied by an attention network linking encoder and decoder; this gives priority to certain vectors, not unlike the way long-term potentiation (LTP) serves to strengthen synapses for better signal transmission between two neurons in the brain.

NMT relies on a large body of parallel corpora, such as *Wikipedia*,²⁶⁾ which are widely available for Japanese-English. Unlike Maori, for instance, Japanese is not a low-resource language in terms of such bitext datasets, but there may still be some specialist fields where insufficient data means that AI systems need the support of human editors and checkers. Xuedong Huang, a technical fellow at Microsoft Research, explains how the latest translation systems, which have been trained using a wealth of documents, learn “the transformation of text”:

Rather than writing handcrafted rules to translate between languages, modern translation systems approach translation as a problem of learning the transformation of text between languages from existing human translations and leveraging recent advances in applied statistics and machine learning.²⁷⁾

There is some similarity between how AI systems perform image recognition and translation (the transformation of text). An AI system that has been given a large dataset consisting of cat and dog pictures will be able to identify pets in new photographs without being given any rules or descriptions: “The key to [machine] learning ... is not *understanding*

but *experience*.”²⁸⁾

Google Translate is certainly far from perfect²⁹⁾ – particularly when working with long, convoluted Japanese sentences. However, failures are far fewer and, depending on the content, the result can sometimes rival the output of weaker students. Whether this serves to encourage or discourage is uncertain, but it should be pointed out that, although computers can now best humans at the most complex board games such as Japanese chess, it will be a long time before they can hope to rival the best human translators at transforming complex Japanese text into English. This is because it requires close reading and context-aware comprehension. The translator has to fully digest the content, nuances, tone, etc. of the source material in order to render it in English. Simply swapping words or phrases may be quite successful when translating between most European languages (except Hungarian and Finnish). This approach may even produce acceptable results when working with fresh iterations of boilerplate text, with which the translation system has been supplied. But to render Japanese prose into good English prose requires not just translation but also editing.

Editing here means the addition, removal and reordering of material. All three of these processes require a clear understanding of the source material, its context and readership. AI will have to improve by leaps and bounds before it can accomplish this. Let us look at some examples.

First, the most common instance in which material must be added when translating from Japanese to English is when the subject is not stated. Here, humans have the distinct advantage over machines. As seen above, phrase-based software typically introduces “I” or “it” whenever it detects a missing subject; this often leads to considerable confusion. In June 2010, Reverso turned 「せっかくの見える化も競争力にはならない」 (“Such investments in *mieruka* will not improve your competitive edge.”) into “Precious, I

see, becoming it don't become a competitive power.” For comparison, Google Translate, without the benefit of any context, will now produce “Visualization of precision does not become competitive” – a clear improvement, but still wanting. No Japanese student would have misunderstood せっかくの見える化.

A common instance in which material in the source text must be removed is repetition, which is not uncommon in written Japanese. That repetition is avoided in English, except for emphasis, can be taught to a computer, but if identical words are not used, how is the computer to spot the repetition?

Japanese prose may also include text or punctuation that is, in English, too obvious to be worth stating (as in “human feelings”), incongruous (as in “approximately 33.1 mm”), or simply unwanted (the ubiquitous hook brackets 「and」). Single hook or corner brackets often serve little discernible function and need to be discarded 3 times out of 4,³⁰ although MT systems will slavishly render these all as quotation marks. In contrast, I advise the students to ignore all single hook brackets in a Japanese text. Then, when polishing the English, quotation marks can be inserted, but only if there is a clear need for them.

Reordering – at the phrase, sentence or even paragraph level – is often called for when translating Japanese to English. Japanese students are initially reluctant to do this, other than inside the sentence, and require encouragement. How, though, can a machine translation system be expected to learn this? This may be the greatest challenge for MT engineers.

Evaluation

The title of a recent article by Sharon Zhou poses and answers a question that people without knowledge or experience of translating

may well ask: “Has AI surpassed humans at translation? Not even close!”³¹⁾ Zhou explains NMT³²⁾ and describes the three main flaws of this approach to machine translation: reliability, memory, and common sense. “NMT is in fact still flawed in many ways that are distinctively _un_human_,” she emphasizes.

The first flaw is reliability. AI translation systems will have to be reliable in future, but because of the past decade of poor machine translations, people will be reluctant to rely on them. NMT systems “often miss negation, whole words, or entire phrases,” Zhou writes.

The second flaw is “really acute short-term memory loss” because current NMT systems translate a sentence at a time. While this is a significant improvement on phrase-based systems, information is not transferred from previous sentences.

The third flaw is that NMT systems have very little common sense, such as “external context and knowledge about the world”.

Zhou also warns against relying on such metrics of quality as the BLEU score, which can lead to such bold claims as “Google’s AI translation system is approaching human-level accuracy”.³³⁾ Unfortunately, it is likely that many more people have noticed such eye-catching titles than have read Zhou’s article. They might wonder if there is any point in studying how to translate, or any future in such a career – especially for language pairs like English & French, Spanish & German.

Zhou’s criticisms are valid, but having seen how far MT has come since I started teaching translation at Keio, I cannot fail to be impressed.

Conclusion

For most people, in the context of JE translation the term “topsy-turvy” would immediately bring to mind the advice that it is often best to start at

the end of a sentence – something which invariably gives rise to problems for the simultaneous interpreter.³⁴⁾ It may also be applied to the challenge of translating out of one’s mother tongue, as is the case with many of the students taking my course. But it also reflects the practice of humans learning to translate from machines. The latest deep learning system behind Google Translate draws on millions of examples to render languages into and out of English. Japanese students, in contrast, have far less experience, but they can rely on much more common sense and contextual awareness than is yet available to the machines. Repeatedly referring to Google Translate output in my course – highlighting both what is impressive and what is plainly wrong – is, in a sense, treating Google Translate as a virtual class member. Whereas machine translations of the past were almost universally poor, Google Translate allows me to set the bar higher: Japanese students have a good chance of beating the machine as they try to create model translations of their own, in which the output is superior to the input. Of course, Google is continuing to improve AI and the bar will rise even higher,³⁵⁾ but for the foreseeable future we can expect Google Translate to remain a worthy competitor in the classroom.

Notes

- 1) See Emma Boyle, “Google Translate Update Makes It Pretty Much as Good as a Human Translator,” *TechRadar*, 28 September 2016.
- 2) See Robin D. Gill’s *Topsy-turvy 1585: A Translation and Explication of Luis Frois S.J.’s Tratado (treatise) Listing 611 Ways Europeans & Japanese are Contrary* (Paraverse Press, 2004). Luis Frois, a Jesuit who lived in Japan, comprehensively listed the ways in which he saw Europeans and Japanese as being contrary to one another – we sniff the top of our melons to see if they are ripe, while they sniff the bottom of theirs, for example.
- 3) “Yoshiaki Nakamura: The Craft of Carpentry,” aired by the BBC on 23 June 2017.

See www.bbc.co.uk/programmes/p056qcfh.

- 4) Jim Breen's JMdict/EDICT Project dates back to 1991.
- 5) Note also that there should be a degree of flexibility in a TM system so that matches do not necessarily have to be identical.
- 6) In other words, read and digest the original, but do not look at it when translating.
- 7) Jim Breen's Japanese-English Dictionary file: www.edrdg.org/jmdict/edict.html.
- 8) There is, however, at least one smartphone application (for both major operating systems) that offers similar functionality: EBPocket.
- 9) The correct translation is "May your business be increasingly prosperous."
- 10) Arthur Waley is not thought to have translated directly from a manuscript.
- 11) Basil H. Chamberlain, *The Kojiki: Records of Ancient Matters* (C.E. Tuttle, 1982) p. 20. This opens the famous passage in which the deities Izanagi and Izanami procreate.
- 12) A useful font for this is CC Wild Words Roman (fontsgoek.com/fonts/CC-Wild-Words-Roman).
- 13) See Sean Leonard, "Celebrating Two Decades of Unlawful Progress: Fan Distribution, Proselytization Commons, and the Explosive Growth of Japanese Animation," *UCLA Entertainment Law Review*, Spring 2005, vol. 12, no. 2.
- 14) Developed by Niels Martin Hansen and Rodrigo Braz Monteiro, this open-source program is remarkably sophisticated and cross-platform. The latest version is 3.2.2 (7 December 2014); it can be downloaded for Microsoft Windows and Mac OS X from www.aegisub.org/downloads.
- 15) One of the better tutorials (youtu.be/z7E6IyUY9ik) has been produced as part of a translation project by a foreign student of Japanese. The method it describes helps the student to focus more on the translation and less on the timings.
- 16) Despite the fondness shown by Japanese media for the subtitling of TV programs, it has proved virtually impossible to obtain Japanese movies with Japanese subtitles. This would seem to show little consideration for the deaf on the part of the distributors. If Kurosawa's *Rashōmon* (1950), for example, were furnished with subtitles, scenes could be used in a JE translation course. Japanese subtitles are, of course, provided for all English-language movies in Japan, so if the sound is muted, these can be used for practice and afterwards a student's translation can be

compared with the original soundtrack.

- 17) VLC media player (www.videolan.org/vlc/) is used for playback of the source with the addition of individual student subtitles, selected from a drop-down menu. It provides fine control of playback and thus facilitates the feedback process.
- 18) Perhaps the only feature I would like added to this software is the ability to create incremental subtitles in which sequential components of a single subtitle appear individually to match pauses in the original dialog.
- 19) Soft-subbing refers to providing subtitles as a separate text file. Using special software, the subtitles can later be “burned” into the source video file (hard-subbing), after which there is no option to switch them off.
- 20) For example, コースもあるようですが, カルトでお願いしました (in a restaurant review) is translated by Google Translate as “Although there seems to be a course, I asked by a cult.” This is then translated back to コースがあるようですが, 私はカルトによって尋ねました. False friends in *katakana* can puzzle AI as well.
- 21) The Bilingual Evaluation Understudy (BLEU) Score is calculated by comparing a generated (in this case, machine-translated) text with a good-quality reference text (in this case, produced by a skilled human translator using the same source material). The result will range from 1.0 (perfect match) to 0.0 (perfect mismatch).
- 22) Japanese was included in the first group of 8 languages, other than English, supported from 2016.
- 23) Until 2007, for languages other than Arabic, Chinese and Russian, Google Translate relied on the technology of SYSTRAN, a machine translation company founded in 1968.
- 24) English thus plays an interlingual role, for now. Although Google Translate will be able to translate directly between language pairs that do not include English, such “zero-shot” translations have yet to be implemented.
- 25) Jeremy Hsu, “Google Translate Gets a Deep-Learning Upgrade,” *IEEE Spectrum*, 3 October 2016.
- 26) Japanese students are frequently surprised when it is pointed out that either there is no Japanese version of a *Wikipedia* page, or that the Japanese is merely a very condensed summary of the English. See, for example, “Fabula and syuzhet” (en.

- wikipedia.org/wiki/Fabula_and_syuzhet). Conversely, the number of *Wikipedia* pages in Japanese only, such as for 日本教 and こけら落とし, appear to be very few.
- 27) As quoted by Emma Woollacott, “How Translation Apps Iron Out Embarrassing Gaffes,” *BBC News*, 25 September 2018.
 - 28) W. P. McNeill, “The Math of Cats and Dogs: AI in Image Recognition,” 27 August 2018.
 - 29) This is perhaps most entertainingly demonstrated by the Translator Fails series of YouTube parodies of famous song lyrics – e.g. Google Translate Sings “Hello” by Adele.
 - 30) While text inside so-called double hook brackets (or white corner brackets) needs to be either rendered in italics or placed between quotation marks, modern Japanese writers tend to use the standard (black) corner brackets to set off a word or phrase from the surrounding text. The human translator might therefore use underlining, a bold font, or highlighting to reproduce the intended effect, though the average English reader would perceive such special treatment as unnecessary.
 - 31) Sharon Zhou, “Has AI Surpassed Humans at Translation? Not Even Close!” *SkyNet Today*, 25 July 2018.
 - 32) For this paper, as in the Zhou’s article, “AI translation” is another name for NMT; the terms are used interchangeably.
 - 33) Nick Statt, “Google’s AI Translation System Is Approaching Human-level Accuracy,” *The Verge*, 27 September 2016; Hany Hassan Awadalla et al., “Achieving Human Parity on Automatic Chinese to English News Translation,” v2, 29 June 2018.
 - 34) This is especially true if the speaker ends a statement with と考えていません (I do not think...) or similar. Translators of text are free from such problems, but on the other hand are usually unable to ask for clarification of what appears ambiguous.
 - 35) Bernard Marr, “Will Machine Learning AI Make Human Translators an Endangered Species?” *Forbes*, 24 August 2018.

Works Cited

- Awadalla, Hany Hassan et al. "Achieving Human Parity on Automatic Chinese to English News Translation." v2, 29 June 2018, arxiv.org/pdf/1803.05567.pdf. Accessed on 24 July 2018.
- Boyle, Emma. "Google Translate Update Makes It Pretty Much as Good as a Human Translator." *TechRadar*, 28 September 2016, www.techradar.com/news/computing/google-translate-update-makes-it-pretty-much-as-good-as-a-human-translator-1329507. Accessed 1 November 2018.
- Chamberlain, Basil H. *The Kojiki: Records of Ancient Matters*. 1892. C.E. Tuttle, 1982.
- Gill, Robin D. *Topsy-turvy 1585: A Translation and Explication of Luis Frois S.J.'s Tratado (treatise) Listing 611 Ways Europeans & Japanese are Contrary*. Paraverse Press, 2004.
- Hsu, Jeremy. "Google Translate Gets a Deep-Learning Upgrade." *IEEE Spectrum*, 3 October 2016, spectrum.ieee.org/tech-talk/computing/software/google-translate-gets-a-deep-learning-upgrade. Accessed 12 November 2018.
- "Japanese Punctuation," *Japan Reference*, www.jref.com/articles/japanese-punctuation.52/. Accessed 19 August 2018.
- Leonard, Sean. "Celebrating Two Decades of Unlawful Progress: Fan Distribution, Proselytization Commons, and the Explosive Growth of Japanese Animation." *UCLA Entertainment Law Review*, Spring 2005, vol. 12, no. 2. escholarship.org/uc/item/0n29d9kq. Accessed 9 December 2018.
- Marr, Benard. "Will Machine Learning AI Make Human Translators an Endangered Species?" *Forbes*, 24 August 2018. www.forbes.com/sites/bernardmarr/2018/08/24/will-machine-learning-ai-make-human-translators-an-endangered-species/#cb06f4a39023. Accessed 9 September 2018.
- McNeill, W. P. "The Math of Cats and Dogs: AI in Image Recognition." 27 August 2018, www.sparkcognition.com/2018/08/27/math-about-cats-dogs/. Accessed 1 September 2018.
- Statt, Nick. "Google's AI Translation System Is Approaching Human-level Accuracy." *The Verge*, 27 September 2016, www.theverge.com/2016/9/27/13078138/google-translate-ai-machine-learning-gnmt. Accessed 21 July 2018.

Woollacott, Emma. "How Translation Apps Iron Out Embarrassing Gaffes." *BBC News*, 25 September 2018, www.bbc.com/news/business-45624685. Accessed 3 October 2018.

"Yoshiaki Nakamura: The Craft of Carpentry." *The Art of Japanese Life*, BBC, 23 June 2017. See www.bbc.co.uk/programmes/p056qcfh.

Zhou, Sharon. "Has AI surpassed humans at translation? Not even close!" *SkyNet Today*, 25 July 2018, www.skynettoday.com/editorials/state_of_nmt. Accessed 12 September 2018.