Gushing Prose
Will Machines Ever be Able to Translate as Badly as Humans?

Rossa Ó Muireartaigh[*]

Flow without Flaw

Type in the word "bowl" into a software translation program, such as Google Translate, and watch as the equivalent Japanese "ボウル" appears almost the moment you press the keys. It is a translation provided instantly, flawlessly, without any of the tortured hesitations and indecision a human translator would experience. Hand the same text "bowl" to the human and watch the paroxysm of self-doubt as they fluster and fret, maybe even grab their heads, and wrack their brains over whether to say "ボウル", "茶碗" or something else. When finished, the human translator will never really know, heart of hearts, whether their translation was really the best one possible, or whether they mistakenly chose the less best option. The machine has done what was asked of it with the resources it was given. It has made no mistake. When the machine gives us an answer we do not want, (such as when I type in "bowl" to mean throwing a ball at a wicket but get uncricket "ボウル") the machine has not made a mistake. It did what I asked. My question was simply wrong. If I build a watermill but flow in sludge rather than water and block the wheel, it was not the watermill's fault, but my own for feeding it the wrong inputs. Not so human translation which is like a watermill that will stop and pause and groan and reverse and spin, once fast and then slow, creaking and squeaking with panic, even when fed a perfect flow of water.

The Wheels of Cognition

What is the difference between a machine and a human translating? Imagine we could watch the process as an outside observer, the same way we can watch water pushing a watermill to grind corn, with each step in the process following

[*] Associate Professor, Department of British and American Studies, School of Foreign Studies, Aichi Prefectural University, 1522-3 Ibaragabasama, Nagakute-shi, Aichi, 480-1198 Japan. Email: omuireartaigh[a]gmail.com

the next. We see the inputted text (the source text) flowing in. Its arrival pushes out a target text. In between, what has happened is that the source text has been converted into some kind of code which is matched with equivalent codes from which the target text will be assembled. There is a neat inevitable flow, each effect has had its singular cause, as the whole translating machine cranks into a symphonic whirl of texts and codes, one pushing the other as the target text is cranked out with all the mechanical reliability of cogs turning cogs. The whole process is automatic. One step producing the next step with the same blind will by which in a watermill water pushes one flap making it rise slightly, to present another flap for water to push, launching the mill into full rotation to turn the connected shaft, that will drive the runner stone to grind the wheat. Nothing in the water or the mill or the runner stone actually knows what it is doing. It is simply parts moving other parts according to the laws of physics. So too with machine translation. Nothing in the machine or the source text or the target text knows what is going on. It is all the action of no-mind, one could say, like a zombified samurai swishing his sword when doing his bushido thing. This is what John Searle's famous thought experiment, entitled the "Chinese Room", so deftly illustrates.[1] Let us imagine someone sitting in a room whose job it is to look at cards with certain codes written on them and to output corresponding cards according to an instruction manual detailing which output card goes with which input card. The cards show Chinese characters which the person reading the manual does not understand. Even so, if the instruction manual is correct and the job properly done, anyone inputting cards into the room would be under the impression that the Room is communicating with them and has, as Alan Turing would argue, a conscious mind.[2] And yet, as Searle's description shows, this is not the case. By pure unconscious mechanical operations language has been used. My point here, of course, being that machine translation is working the same way, outputting in response to what is inputted without ever knowing what it is doing. It is as flawless and devoid of mistakes in its proper operation as a watermill grinding wheat, but equally unknowing of what it is doing.

Blackbox Beetles

With humans it could not be more different. For no matter how much we try

---

[1] John Searle, "Minds, brains, and programs," *Behavioral and Brain Sciences* 3 (3) (1980): 417-457.

[2] Alan Turing, "Computing Machinery and Intelligence," *Mind* 59 (236) (1950): 433–60.

to observe each step of the translating operation when done by a human, there will be a certain blackbox moment where we cannot peer in any further. This is the moment when the translator is consciously choosing which words to use, or to borrow Searle's image again, which cards to output. For in the case of the human, there is no instruction manual, or rather the instruction manual lies locked and hidden within the blackbox of the conscious mind. Maybe it is not even there at all. This blackbox I speak of can be thought of as the beetlebox in yet another famous thought experiment, this time by Wittgenstein.[3] The idea of the beetlebox is to imagine that I hold in my hands a box which may or may not contain a beetle. You can never open the box to find out. And yet if I keep using the word "beetle" correctly in our shared language, the box might as well have a beetle in it. In other words, when we use language with another person we can only judge what is going on in their minds by how they use language. This secondary layer of evidence, the visible, surface presence of language, reflects, we can only assume, an inner consciousness. It does not grant us full direct sight of it. Let us take, for example, the word '茶碗' [chawan]. When I use this word you have absolutely no idea if my understanding of it is the same as yours. So long as I continue to use the word correctly, to play the language game correctly (or let's call it "translation game"), such as when I utter the sentence "this *chawan* evokes the asesthetics of *sabi-wabi*", you cannot accuse me of not understanding what "*chawan*" means. (On the contrary, if I say, for example, "my *chawan* is up by 5 degrees since I went jogging this morning", you are entitled to wonder if our meanings are the same.) Which means then that if I translate "茶碗" as "cup" you cannot accuse me of not understanding the word, even if you think "*chawan*" or "tea bowl" is the better translation. When we dispute the best possible translations for a source text we are still not looking into each other's blackbox of inner consciousness, we are merely negotiating and contending moves in the translation game. There is no telepathy, only translation games. As an aside, may I mention the fact that I have in my time met orientalists who do, implicitly, claim to have telepathic powers and to be able to look into the inner conscious minds of non-Japanese and non-Asians and convince themselves that the understandings such people have of words such as "茶碗" [chawan], "さび" [sabi] and "わび" [wabi] are not the same as theirs. They have a remarkable ability to see the ignorance in the minds of others, but, alas, not in their own.

---

[3] Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1958).

Seeing Consciousness

The essential difference between the machine translator and the human one is the existence of this consciousness at the core of the human mind—which cannot be seen or accessed. In other words, from the third person point of view, the human translator and the machine translator look exactly the same. The only giveaway, maybe, is the fact that the human translator isn't as good. I can only assume the human translator has a consciousness on the grounds that I have one myself, and because the human is behaving as though they had a consciousness, something I judge on the basis of whether or not they replicate my own behavior. When I translate, I do not translate with the automaticity of a machine, as much as I may try, but instead experience constant indecisions, aporias, and dissatisfaction as I, with conscious deliberation, abduce interpretations of the source text. Other human translators seem, on the surface, to be doing the same thing. On these grounds alone, I assume that they are conscious (and the translating machine is not).

To express this further let me reach for another famous thought experiment, David Chalmers's zombies. [4] Here we are asked to consider the difference between conscious humans and zombies (who can perfectly replicate the behavior of said conscious humans). It would seem that for third party observers telling this difference would be impossible. A zombie that behaved perfectly like a human would be, for us, a human. And we would falsely assume conscious experience where there was none. Similarly, a machine translation could, theoretically, replicate perfectly a human translator, perhaps by being programmed to behave with slow self-torturing hesitancy, shifting between flashes of creative brilliance and dull ploding, all the while demonstrating stressed-out behavior, worrying about impossible deadlines imposed by narky underpaying translation agencies spouting sanctimonious gunk about 'quality control'. In such a case we would then falsely ascribe conscious experience to the machine, apropos again Turing's *gedanken*. Both zombies and translating machines lack something that cannot be seen. And that is, conscious experience. And this can pose a problem for the consciousness-seeking outside observer who ends up having to try to prove a negative, looking for what is not meant to be seen.

---

[4] David Chalmers, *The Conscious Mind* (New York; Oxford: Oxford University Press, 1996).

All or Nothing

Slavoj Žižek in *Organs without Bodies* refers to Chalmers and makes the argument that trying to attain a third-person perspective on consciousness's place in the universe is both an epistemological and ontological impossibility (it isn't to be known and it is not there). Žižek summarizes his view in a paragraph that is worth reading in full:

> When Chalmers writes in his argument against the reductive explanation of consciousness that "even if we knew very last detail about the physics of the universe—the configuration, causation, and evolution among all the fields and particles in the spatio-temporal manifold—*that* information would not lead us to postulate the existence of conscious experience," he commits the standard Kantian mistake: such a total knowledge is strictly nonsensical, epistemologically *and* ontologically. His reasoning is the obverse of the vulgar determinist notion articulated by in Marxism by Bukharin, who wrote that, if we were to know the entirety of physical reality, we would also be able to predict precisely the emergence of a revolution. More generally, this line of reasoning—consciousness as an excess/surplus over physical totality—is misleading since it has to evoke a meaningless hyperbole. When we imagine the Whole of reality, there is no longer any place for consciousness (and subjectivity). There are, as we have already seen, only two options left open here: either subjectivity is an illusion or reality is *in itself* (not only epistemologically) not-All.[5]

Žižek, of course, is the great *pratyekabuddha* of our times. And so, his views dovetail nicely with that of much of the Kyoto School.[6] Reality can never be seen

---

[5] Slavoj Žižek, *Organs without Bodies* (New York; London: Routledge, 2004), 115.

[6] Nishida Kitaro clearly places links world to consciousness in the now: "The world thus begins in the dimension that the present determines itself as the self-determination of the absolute present. The self begins in the dimension that knowledge and will unite, in the sense that that which reflects and that which is reflected are one. The self arises in the place that the world arises; the world arises in the place that the self arises. Therefore, our self-consciousness is the self-expression of the world, and the self-expression of the world is our self-consciousness." ("Toward a Philosophy of Religion with the Concept of Preestablished Harmony as a Guide," *The Eastern Buddhist* 3(1) (1970), 36-37) This placing of consciousness in the now from the which the world is emergent leads to the assumption of an emptiness where the All that can only be the not-All can be. As Nishitani Keiji describes it: "In short, it is only on a field where the being of all things is a being at one with emptiness that it is possible for all things to

in a totality that could be observed by an omniscient third party and the proof of this impossibility is the unplaceability of consciousness when one does adopt such a view. This is an opinion that Nishida Kitaro and Nishitani Keiji, along with Žižek, judging from his take on Chalmers, would hold to. Why is full knowledge of everything not possible? To have full knowledge of everything would be the crudest of pantheism. It would mean that the thoughts in your head would be the same as the thoughts of God. There could be no separation between you and God. And yet you experience yourself as being separate to God (or at least separate to the rest of the world that is not you). This experience, which by definition means that we have a *partial* knowledge of things, would have to be included in the *full* knowledge of things. Not knowing everything would be part of the absolute condition of knowing all. It is an impossibility. Your consciousness, with its partial knowledge, is a constant knowledge gap within the All of existence, of the world, of the universe.

Instead we need to lose the idea that knowledge is the mapping of all that is there onto all that we can consciously make out to be all that is there. Instead of seeing the world out there as a stable body of knowledge on which we flash the torches of our inner consciousness, attaining knowledge of one piece of a giant picture that someone with a bigger flash lamp (like God or a cosmic machine) could see overall, we should see the world out there as not prior known. The world out there, in fact, shapes itself into a stable body of knowledge to suit the flash lamp we shine upon it. Otherwise, without our conscious attention, it is not, in effect, "there". It is empty.

In this picture, knowledge must be relocated at the point of consciousness on the field of emptiness. Only a consciousness, such as that of a human, can be there at such a point. A translation machine, crunching away the bits and signals and codes of a fed-in text is not there. It can only be put there when it is the object of our attention. In other words, when a machine translates, it is because we have let ourselves believe that it has translated. The machine translates and the watermill turns because our minds translate and our minds turn.

The Mystery of Translation

To consider the special place (or non-place) of the conscious mind, let us

---

gather into one, even while each retains its reality as an absolutely unique being." (*Religion and Nothingness*, 1982, 148).

consider further how the act of translation actually works. When we translate we produce a new text that does not look in anyway like the old text. And yet there is a sameness about them, otherwise they would not be translations, but merely two different texts in two different languages. To illustrate, let us take the following Japanese source text:

[source text] 青磁輪茶碗[7]

This can be rendered into English as:

[target text] Celadon bowl with foliated rim.

Other than being black lines on a white page, there is absolutely nothing in the source text that in any way resembles the target text to hint at the fact that they are the same thing. They are no more connected than the following three lines "///" are connected to the word "flower".

And yet a translation has occurred. I, the translator (with the help of an uncredited translator of a museum brochure), "saw" meaning in the source text and was able to create a target text with the same meaning. Where am I finding the meaning that links the two texts? It is not in the surface, visible depictions. Where can it be? It is always tempting to place meaning in translation as residing in-between the source and target texts. We can make this argument in two ways: by focusing on how the world is divided into universal ideas that can then be represented by words in any language (the semantic path), or by focusing on how each language is carved from a finite set of parametric choices that are common to all languages (the syntactic path). Indeed, if we were to build a translating machine, we would look to access this in-between language zone and automate the process by which source and target texts move between this zone. Following my distinction above, we could create two possible programs, a semantic one and a syntactic one.

---

[7] *The Lee Kilsoo Collection—Ceramics of Korea, Japan, China and Southeast Asia/*委吉秀コレクション―韓国・日本・中国・東南アジアのやきもの (Aichi Prefectural Ceramic Museum, 2008), 12. The catalog has rendered it as "Bowl with foliated rim, caledon". I moved "caledon" to the front because I feel that it can be used as a modifying adjective. I grieve that my knowledge will never be enough to know if this was the correct thing to do.

A Semantic Machine

Let us look at the semantic program first. We know, from our everyday experience, that the world is divided into different objects. Only in the realm of magic and fantasy would this not be the case. Our task, then, in programming our machine would be to identify each of these objects. For instance, let us take a small round handle-less piece of porcelain as one such object. I could, let's imagine, give it the label "1.1". I can then attach different labels to Object 1.1., such as "bowl" in English. In fact, I could use words in any other language I wanted (for example, "*mias*" in Irish, "*bovlo*" in Esperanto, or "*bol*" in French). And so when I feed the word "bowl" into my machine, it will automatically connect it to "1.1" and, if doing English to Japanese, will output "鉢", providing a successful translation. The problem, though, with such perfect in-between languages is that ideas do not attach themselves to the world of objects so easily.[8] Or rather, the problem is that the objects of the world do not attach themselves so easily to our ideas.

I can see before me, undeniably and unmistakably, an object that I know to be, without any controversy, a "bowl". It is a coherent universally recognizable physical object transcendent of all cultural perspectives. One could pick it up and throw it to or at anyone who claimed that my seeing it was merely a product of western patriarchal socio-cultural discourses of power. And yet the whole concept of "bowl" starts to get fuzzy and unclear as soon as I start to relate it to other related objects. When is a bowl not a bowl? When it is a basin? Is it defined by shape? If I make a bowl from very thin paper is it still a bowl, even if it cannot really function as a bowl? If a bowl is small and I drink tea out of it, is it still a bowl? Or is it a type of cup?

It seems, then, that the concept of "bowl" is not as solid and indisputable as the hard ceramic object sitting before me but is highly dependent on the particular local context in which it is being used. A *bowl* can become a *basin* or *cup* or a *paper craft* depending on the situation in which we are using the concept. Indeed, the more common the object, the more slippery the concept. The irony is that the more obscure and domain-specific a word is the easier it is to translate because the less chance there is of the concept having slid naturally into other concepts. For instance, when I input the words "青磁輪茶碗" into the machine translation

---

8 See Eco, *The Search for the Perfect Language* (Oxford; London: Blackwell, 1995), for a fuller history of the centuries-old unsuccessful quest for a perfect pure in-between language.

program SYSTRAN (an historic leader in the field) I get "celadon tea bowl". (Which is very impressive: Google Translate gave me "celadon wheel tea bowl"). However, when I type in simply "茶碗" I get the rarified "*chawan*." But if truth be told, if I were to pick up a 茶碗 and speak of it in English, I would say "cup", "little cup", "saucer", "porcelain piece," "tea cup," "tea bowl," "thing," "this thing," "small mug," "yoke," "tankard type thing" and a dozen other words. However, I don't think I would ever say "*chawan*." (I lack the *savoir faire* for italicized words). Could a computer ever be programmed to replicate the messiness, inconsistencies, vagueness, and borderline inarticulateness of my language usage? Universal interlanguage fails because we humans are too confusing in our descriptions of the world we live in and too unpredictable in how we go about dividing it and categorizing it. We never really are knowing what we are going to say and make it up as we go along. Unlike machines, which often seem to have already made up what we never even knew we didn't want to say. Meaning is local and hinged to present usage. For a language to work as a language it does have to present itself as being absolute in its semantics (a bowl is a bowl) but this absoluteness works only in so far as the contingent use of language allows it. Machines can never get at these contingencies, and processing is always handicapped and limited by the machine's dependence on absolute semantics. To do meaning one must be conscious in a present place where all the world converges in one's interpretation of the now moment. With this, it can be said that, quite simply, machines don't do meaning.[9]

A Syntactic Machine

Another way to build our translation machine would be to separate the meaning from the syntax (so the world out there, the semantics, is not a problem) and just let computers work with and replicate the universal parameters from

---

[9] This seems to be the argument Hurbert Dreyfus is making in his book *What Computers Can't Do: A Critique of Artificial Reason* (New York: Harper and Row, 1972), 200: "A phenomenological description of our experience of being-in-a-situation suggests that we are always already in a context or situation which we carry over form the immediate past and update in terms of events that in the light of this past situation are seen to be significant. We never encounter meaningless bits in terms of which we have to identify the context, but only facts which are already interpreted and which reciprocally define the situation we are in. Human experience is only intelligible when organized in terms of a situation in which relevance and significance are already given. This need for prior organization reappears in AI as the need for a hierarchy of contexts in which a higher or broader context is used to determine the relevance and significance of elements in a narrower or lower context."

which, as the Chomskyanists tell us (and nobody has as yet proved them wrong) all languages are carved. A sentence in a language is arranged according to an underlying syntactic structure which is systematic and runs regardless of the words being fed into it. "Colorless celadon bowls with foliated rim sleep furiously" is semantically meaningless but grammatically correct because it follows the syntactic patterns of English. Every language has its own syntactic patterns but these are all shaped from a common universal set of parameters which are either activated or not activated according to the dictates of the particular language. At first glance, then, it may seem that this could offer a possibility for machine translation since we now have an interlanguage that is systematic and not messed up by human fuzzy logic. Removing meaning from the equation means that meaning-illiterate machines are now mean and lean and ready to play. Feed in the sentence "a bowl fell". The machine slices it into a Noun Phrase ("a bowl") and a Verb Phrase ("fell"), sorts it and outputs "不定の単数のカップが落ちた". No problem. However, the limit we find with this syntactic interlanguage is that each language, as it switches on and off parameters, finds itself committed to providing certain information in its grammar that other languages do not have to give. Thus, for instance, English, because of the parameters it has inherited, is a language where singular and plural must *always* be designated. This is not the case in Japanese, for example. There is no problem with this when we go from English to Japanese. "A bowl" or "bowls" can be simply "ボウル". But it is a problem when we go the other way. If a text consists of the Japanese word "ボウル", the English translator does have to decide if it is one or more bowls. In other words, the further we go down the branching parameters the harder it is for us to go back up again, making translation, without semantic knowledge, an impossibility. Chomskyan linguistics does admirably explain how all languages are chiseled out of the same block. However, it does have problems explaining how translation can pass, horizontally, between the vertical branches of a sentence's underlying syntactic structures.

It seems, then, that any machine translation that seeks to translate as humans do through the search for a pure in-between set of concepts or grammatical patternings is doomed to failure. When a text is cut into bits this actually changes each bit and they can never be put back together again. This is the mystery of translation. It happens when both the source and target texts remain completely intact. This is why the most successful automated translation programs, as measured by the erratic and irrational expectations of humans, are those that use

translations that have already been done by humans and basically copy them when other similar source texts arise.[10] In other words, computers can translate like humans only when they cheat and plagiarize those humans.

Form is Emptiness

The fact that human translation occurs even though there is no in-between meaning or common structure between languages needs to be fully grasped. Let us ask again the question, between two texts that are translations of each other where does the meaning lie? If not in-between then where? We cannot say in both, as this is simply restating the problem. For meaning to be in both, they both must have something in common. But that something in common is not, as we have seen, anywhere else but in the two texts, both of which are, in form, completely different from one another. Meditate again on the two texts: "青磁輪茶碗" and "Celadon bowl with foliated rim". See again how different they are. If you say they are the same because they both express the same "concepts" or "ideas" you are, again, trying to locate an in-between realm that does not exist anywhere—except in the two texts that are so different. Ideas and concepts are *empty* without the texts to express them. The only way beyond this impasse is to realize that the form is this emptiness, and this emptiness is the form. In other words, we must remember the words of the *Heart Sutra*: 空不異色、色不異空 ("emptiness differs not from form, form differs not from emptiness"). In other words, when we stare at two translations, so different in form, the only way they can be connected is in our consciousness. But our consciousness can do this because, unlike machines which only ever move between form and form, it can produce meaning by emptying emptiness to produce form. Humans translate because

---

[10] As *The Economist* reports ("Language: Finding a Voice," May 1, 2017): "Many early approaches to language technology—and particularly translation—got stuck in a conceptual cul-de-sac: the rules-based approach. In translation, this meant trying to write rules to analyze the text of a sentence in the language of origin, breaking it down into a sort of abstract "interlanguage" and rebuilding it according to the rules of the target language. These approaches showed early promise. But language is riddled with ambiguities and exceptions, so such systems were hugely complicated and easily broke down when tested on sentences beyond the simple set they had been designed for. Nearly all language technologies began to get a lot better with the application of statistical methods, often called a "brute force" approach. This relies on software scouring vast amounts of data, looking for patterns and learning from precedent. For example, in parsing language (breaking it down into its grammatical components), the software learns from large bodies of text that have already been parsed by humans. It uses what it has learned to make its best guess about a previously unseen text. In machine translation, the software scans millions of words already translated by humans, again looking for patterns."

human consciousness, that radical self-reflectivity that can only be from a point of emptiness, (where the eye cannot see the eye), where meaning, and not preprogrammed mechanical change, is generated. Machines churn out their translations in a flawless flow. They miss nothing. But that is precisely their problem.

Bibliography

Chalmers, D. J. *The Conscious Mind: In Search of a Fundamental Theory*. New York; Oxford: Oxford University Press, 1996.

Dreyfus, Hubert L. *What Computers Can't Do: A Critique of Artificial Reason*. New York: Harper and Row, 1972.

Eco, Umberto. *The Search for the Perfect Language.* Oxford; London: Blackwell, 1995.

Nishida, Kitaro. 1970. "Toward a Philosophy of Religion with the Concept of Preestablished Harmony as a Guide." Translated by David Dilworth. *The Eastern Buddhist*, 3(1): 19–46.

Nishitani, Keiji. *Religion and Nothingness*. Translated by Jan Van Bragt. Berkeley: University of California Press, 1982.

Searle, John. R. "Minds, brains, and programs." *Behavioral and Brain Sciences.* 3(3) (1980): 417-457 [Pre-print version]

Turing, A., "Computing Machinery and Intelligence," *Mind.* 59 (236) (1950): 433–60.

Wittgenstein, Ludwig. *Philosophical Investigations.* Translated by G. E. M. Anscombe. Oxford: Blackwell, 1958.

Žižek, Slavoj. *Organs without Bodies: On Deleuze and Consequences*. New York; London: Routledge, 2004.