The New York Times Magazine

https://www.nytimes.com/2015/06/07/magazine/is-translation-an-art-or-a-math-problem.html

THE TECH & DESIGN ISSUE

Is Translation an Art or a Math Problem?

By Gideon Lewis-Kraus

June 4, 2015

One Enlightenment aspiration that the science-fiction industry has long taken for granted, as a necessary intergalactic conceit, is the universal translator. In a 1967 episode of “Star Trek,” Mr. Spock assembles such a device from spare parts lying around the ship. An elongated chrome cylinder with blinking red-and-green indicator lights, it resembles a retracted light saber; Captain Kirk explains how it works with an off-the-cuff disquisition on the principles of Chomsky’s “universal grammar,” and they walk outside to the desert-island planet of Gamma Canaris N, where they’re being held hostage by an alien. The alien, whom they call The Companion, materializes as a fraction of sparkling cloud. It looks like an orange Christmas tree made of vaporized mortadella. Kirk grips the translator and addresses their kidnapper in a slow, patronizing, put-down-the-gun tone. The all-powerful Companion is astonished.

“My thoughts,” she says with some confusion, “you can hear them.”

The exchange emphasizes the utopian ambition that has long motivated universal translation. The Companion might be an ion fog with coruscating globules of viscera, a cluster of chunky meat-parts suspended in aspic, but once Kirk has established communication, the first thing he does is teach her to understand love. It is a dream that harks back to Genesis, of a common tongue that perfectly maps thought to world. In Scripture, this allowed for a humanity so well coordinated, so alike in its understanding, that all the world’s subcontractors could agree on a time to build a tower to the heavens. Since Babel, though, even the smallest construction projects are plagued by terrible delays.

Translation is possible, and yet we are still bedeviled by conflict. This fallen state of affairs is often attributed to the translators, who must not be doing a properly faithful job. The most succinct expression of this suspicion is "*traduttore, traditore*," a common Italian saying that's really an argument masked as a proverb. It means, literally, "translator, traitor," but even though that is semantically on target, it doesn't match the syllabic harmoniousness of the original, and thus proves the impossibility it asserts.

Translation promises unity but entails betrayal. In his wonderful survey of the history and practice of translation, "Is That a Fish in Your Ear?" the translator David Bellos explains that the very idea of "infidelity" has roots in the Ottoman Empire. The sultans and the members of their court refused to learn the languages of the infidels, so the task of expediting communication with Europe devolved upon a hereditary caste of translators, the Phanariots. They were Greeks with Venetian citizenship residing in Istanbul. European diplomats never liked working with them, because their loyalty was not to the intent of the foreign original but to the sultan's preference. (Ottoman Turkish apparently had no idiom about not killing the messenger, so their work was a matter of life or death.) We retain this lingering association of translation with treachery.

The empire of English has a new Phanariot class, and they are inventing the chrome light-saber apps of the utopian near-future. They are native speakers of C++, and they reside in our midst on semipermanent loan from the Internet. On the plus side, they are faithful to no sultan. The minus is that they are not particularly loyal to any language at all.

Google Translate is far and away the venture that has done the most to realize the old science-fiction dream of serene, unrippled exchange. The search giant has made ubiquitous those little buttons, in email and on websites, that deliver instantaneous conversion between language pairs. Google says the service is used more than a billion times a day worldwide, by more than 500 million people a month. Its mobile app ushers those buttons into the physical world: The camera performs real-time augmented-reality translation of signs or menus in seven languages, and the conversation mode allows for fluent colloquy, mediated by robot

voice, in 32. There are stories of a Congolese woman giving birth in an Irish ambulance with the help of Google Translate and adoptive parents in Mississippi raising a child from rural China.

Since 2009, the White House's policy paper on innovation has included, in its list of near-term priorities, "automatic, highly accurate and real-time translation" to dismantle all barriers to international commerce and cooperation. If that were possible, a variety of local industries would lose the final advantage of their natural camouflage, and centralization — in social networking, the news, science — would accelerate geometrically. Nobody in machine translation thinks that we are anywhere close to that goal; for now, efforts in the discipline are mostly concerned with the dutiful assembly of "cargo trucks" to ferry information across linguistic borders. The hope is that machines might efficiently and cheaply perform the labor of rendering sentences whose informational content is paramount: "This metal is hot," "My mother is in that collapsed house," "Stay away from that snake." Beyond its use in Google Translate, machine translation has been most successfully and widely implemented in the propagation of continent-spanning weather reports or the reproduction in 27 languages of user manuals for appliances. As one researcher told me, "We're great if you're Estonian and your toaster is broken."

Warren Weaver, a founder of the discipline, conceded: "No reasonable person thinks that a machine translation can ever achieve elegance and style. Pushkin need not shudder." The whole enterprise introduces itself in such tones of lab-coat modesty. The less modest assumption behind the aim, though, is that it's possible to separate the informational content of a sentence from its style. Human translators, like poets, might be described as people for whom such a distinction is never clear or obvious. But human translators, today, have virtually nothing to do with the work being done in machine translation. A majority of the leading figures in machine translation have little to no background in linguistics, much less in foreign languages or literatures. Instead, virtually all of them are computer scientists. Their relationship with language is mediated via arm's-length protective gloves through plate-glass walls.

Many of the algorithms used by Google and Skype Translator have been developed and honed by university researchers. In May, a computational linguist named Lane Schwartz, who teaches at the University of Illinois at Urbana-Champaign, hosted the first Machine Translation Marathon in the Americas, a weeklong hackathon to improve the open-source tools that those without Google resources share. Urbana-Champaign is largely known outside Illinois for two people: David Foster Wallace, who grew up there, and Marc Andreessen, who invented the first widely adopted graphical web browser as a student at the university. (Schwartz suggested a third: HAL 9000.) It is tempting to see them as the two ends of a spectrum: Wallace as a partisan of neologism, allusion and depth, Andreessen on the side of proliferation, access and breadth.

At this conference, at least, the spirit of Andreessen prevailed. Though attendees hailed from places like Greece, India, Israel, Suriname and Taiwan, almost nobody betrayed any interest in language as such. They understood that language is a rich and slippery thing, but they were there for the math.

The marathon took place at a conference center attached to something called an iHotel. The center was a U-shaped hallway lined by rooms named after virtues — the Leadership Boardroom, the Loyalty Room, the Knowledge Room, the Innovation Room and the Excellence Room. At the presentations, computer scientists with straight faces regularly made comments like “Paragraphs arguably should be coherent in topic” or “Grammatical structure can matter in a sentence.” One presenter said that sometimes French places its adjective before the noun and sometimes after, but that, he concluded with a short shrug, “nobody knows why or when.”

One of the American marathon presenters wore two consecutive days of threadbare grammar T-shirts — one read, “Good grammar costs nothing!” and the other, “I am silently correcting your grammar” — so I imagined he might see his algorithmic work in the context of broader linguistic interests. I asked him if he spoke any other languages, and he said: “I speak American high-school French,

which is to say I don't. But it's surprising how little it helps to know another language. When you're working with so many languages, it's actually not helpful to know one." (His third T-shirt read, "Don't follow me, I'm lost, too.")

The possibility of machine translation, Schwartz explained, emerged from World War II. Weaver, an American scientist and government administrator, had learned about the work of the British cryptographers who broke the Germans' Enigma code. It occurred to him that cryptographic investigations might solve an immediate postwar problem: keeping abreast of Russian scientific publications. There simply weren't enough translators around, and even if there were, it would require an army of them to stay current with the literature. "When I look at an article in Russian," Weaver wrote, "I say: 'This is really written in English, but it has been coded in some strange symbols. I will now proceed to decode.' " In this view, Russian was merely English in frilly Cyrillic costume, only one small step removed from pig Latin.

Within a year or two, this idea was understood as absurd, and yet the broader notion of algorithmic processing held. By 1954 the American public was treated to a demonstration of the first nonnumerical application of computing. A secretary typed a Russian sentence onto a series of punch cards; the computer whirred and spat out an English equivalent. The Christian Science Monitor wrote that the "electronic brain" at the demonstration "didn't even strain its superlative versatility and flicked out its interpretation with a nonchalant attitude of assumed intellectual achievement."

That demonstration, however, was basically rigged. The computer had been given a pidgin vocabulary (a total of 250 words) and fed a diet of simple declarative sentences. In 1960, one of the earliest researchers in the field, the philosopher and mathematician Yehoshua Bar-Hillel, wrote that no machine translation would ever pass muster without human "post-editing"; he called attention to sentences like "The pen is in the box" and "The box is in the pen." For a translation machine to be successful in such a situation of semantic ambiguity, it would need at hand not only a dictionary but also a "universal encyclopedia." The brightest future for machine translation, he suggested, would rely on coordinated efforts between plodding

machines and well-trained humans. The scientific community largely came to accept this view: Machine translation required the help of trained linguists, who would derive increasingly abstract grammatical rules to distill natural languages down to the sets of formal symbols that machines could manipulate.

This paradigm prevailed until 1988, year zero for modern machine translation, when a team of IBM's speech-recognition researchers presented a new approach. What these computer scientists proposed was that Warren Weaver's insight about cryptography was essentially correct — but that the computers of the time weren't nearly powerful enough to do the job. "Our approach," they wrote, "eschews the use of an intermediate mechanism (language) that would encode the 'meaning' of the source text." All you had to do was load reams of parallel text through a machine and compute the statistical likelihood of matches across languages. If you train a computer on enough material, it will come to understand that 99.9 percent of the time, "the butterfly" in an English text corresponds to "*le papillon*" in a parallel French one. One researcher quipped that his system performed incrementally better each time he fired a linguist. Human collaborators, preoccupied with shades of "meaning," could henceforth be edited out entirely.

Though some researchers still endeavor to train their computers to translate Dante with panache, the brute-force method seems likely to remain ascendant. This statistical strategy, which supports Google Translate and Skype Translator and any other contemporary system, has undergone nearly three decades of steady refinement. The problems of semantic ambiguity have been lessened — by paying pretty much no attention whatsoever to semantics. The English word "bank," to use one frequent example, can mean either "financial institution" or "side of a river," but these are two distinct words in French. When should it be translated as "*banque*," when as "*rive*"? A probabilistic model will have the computer examine a few of the other words nearby. If your sentence elsewhere contains the words "money" or "robbery," the proper translation is probably "*banque*." (This doesn't work in every instance, of course — a machine might still have a hard time with the relatively simple sentence "A Parisian has to have a lot of money to live on the Left Bank.") Furthermore, if you have a good probabilistic model of what standard

sentences in a language do and don't look like, you know that the French equivalent of "The box is in the ink-filled writing implement" is encountered approximately never.

Contemporary emphasis is thus not on finding better ways to reflect the wealth or intricacy of the source language but on using language models to smooth over garbled output. A good metaphor for the act of translation is akin to the attempt to answer the question "What player in basketball corresponds to the quarterback?" Current researchers believe that you don't really need to know much about football to answer this question; you just need to make sure that the people who have been drafted to play basketball understand the game's rules. In other words, knowledge of any given source language — and the universal cultural encyclopedia casually encoded within it — is growing ever more irrelevant.

Many computational linguists continue to claim that, after all, they are interested only in "the gist" and that their duty is to find inexpensive and fast ways of trucking the gist across languages. But they have effectively arrogated to themselves the power to draw a bright line where "the gist" ends and "style" begins. Human translators think it's not so simple. The machinist's attitude is that when someone's mother is trapped under a house, it's fussy and self-important to worry too much about nuance. They see the redundancy and allusiveness of natural languages as a matter not of intricacy but of confusion and inefficiency. Most valuable utterances revert to the mean of statistical probability. If this makes them unpopular with poets and fanciers of language, so be it. "Go to the American Translators Association convention," one marathon attendee told me, "and you'll see — they hate us."

This is to some extent true. As the translator Susan Bernofsky put it to me, "They create the impression that translation is not an art." (A widely admired literary translator, who wished to remain anonymous, admitted that although she worries about machine translation's mission creep, she thinks Google Translate is a wonderful tool for writing notes to the woman who cleans her house.)

What mostly annoys human translators isn't the arrogance of machines but their appropriation of the work of forgotten or anonymous humans. Machine translation necessarily supervenes on previous human effort; otherwise there wouldn't be the parallel corpora that the machines need to do their work. I mentioned to an Israeli graduate student that I had been reading the Wikipedia page of Yehoshua Bar--Hillel and had found out that his granddaughter, Gili, is a minor celebrity in Israel as the translator of the "Harry Potter" books. He hadn't heard of her and didn't seem interested in the process by which a publisher paid to import books about magic for children. But we would have no such tools as Google Translate for the Hebrew-English language pair if Bar-Hillel had not hand-translated, with care, more than 4,000 pages of an extremely useful parallel corpus. In a sense, their machines aren't actually translating; they're just speeding along tracks set down by others. This is the original sin of machine translation: The field would be nowhere without the human translators they seek, however modestly, to supersede.

Perhaps to paper over the associated guilt, the group in Urbana-Champaign cultivated a minor resentment toward their human counterparts. More than once I heard someone at the marathon refer to the fact that human translators are finicky and inconsistent and prone to complaint. Quality control is impossible. As one attendee explained to me, "If you show a translator an unidentified version of *his own translation* of a text from a year ago, he'll look it over and tell you it's terrible."

One computational linguist said, with a knowing leer, that there is a reason we have more than 20 translations in English of "Don Quixote." It must be because nobody ever gets it right. If the translators can't even make up their own minds about what it means to be "faithful" or "accurate," what's the point of worrying too much about it? Let's just get rid of the whole antiquated fidelity concept. All the Sancho Panzas, all the human translators and all the computational linguists are in the same leaky boat, but the machinists are bailing out the water while the humans embroider monograms on the sails.

But like many engineers, the computational linguists are so committed to the power and craftsmanship of their means that they tend to lose perspective on whose ends they are advancing. The problem with human translators, from the time of the Phanariots, is that there is always the possibility that they might be serving the ends of their bosses rather than the intent of the text itself. But at least a human translator asks the very questions — *What purpose is this text designed to serve? What aims are encoded in this language?* — that a machine regards as entirely beside the point.

The problem is that *all* texts have some purpose in mind, and what a good human translator does is pay attention to how the means serve the end — how the "style" exists in relationship to "the gist." The oddity is that belief in the existence of an isolated "gist" often obscures the interests at the heart of translation. Toward the end of the marathon, I asked a participant why he chose to put his computer-science background to the service of translation. He mentioned, as many of them did, a desire to develop tools that would be helpful in earthquakes or war. Beyond that, he said, he hoped to help ameliorate the time lag in the proliferation of international news. I asked him what he meant.

"There was, for example, a huge delay with the Germanwings crash."

It wasn't the example I was expecting. "But what was that delay, like 10 or 15 minutes?"

He cocked his head. "That's a huge delay if you're a trader."

I didn't say anything informational in words, but my body or face must have communicated a response the engineer mistranslated as ignorance. "It's called cross-lingual arbitrage. If there's a mine collapse in Spanish, you want to make a trade as quickly as possible."

Gideon Lewis-Kraus is a contributing writer for the magazine and the author of "A Sense of Direction: Pilgrimage for the Restless and the Hopeful."

Sign up for our newsletter to get the best of The New York Times Magazine delivered to your inbox every week.

A version of this article appears in print on , Page 48 of the Sunday Magazine with the headline: Tower of Babble