

Slave to the Word

By Michael Karpa

Michael Karpa is a long-time Japanese-to-English translator based in San Francisco, California. He has a master's degree in international policy from Stanford University and was a visiting professor at the Monterey Institute of International Studies Graduate School of Translation and Interpretation in 2002-2003. He is the translator of Sakaiya Taichi's What Is Japan and his fiction and essays have been published in Faultline and a number of other literary magazines.

TRANSLATORS—human translators—have often greeted the prospect of computer translation with naked fear. Will we be out-competed by efficient machines, as so many professions have? We love the productivity-boosting Web, home of online dictionaries, search engines, and even translation memory tools, but Web-based statistical machine translation, which essentially catalogs the translation labors of the past, is making tremendous strides. Perhaps our days are numbered.

But not so fast. Against this mechanical efficiency, human translators offer “understanding.” In a recent article, I speculated that the way translation forces us to demonstrate the completeness of our understanding could instead increase the value of what we do.¹ A growing body of evidence suggests that intensive exposure to the Web is impairing our cognitive abilities: to read, to understand, to think.² The general idea is that the decision-making the Web demands (whether to click links or not) overloads our working memory, the part of the brain that holds the information we are using at the moment, preventing it from forming those links to other information, experiences, and emotions that occur when we engage in creative thought. © Translators are practiced in drawing constantly on inference, knowledge, and

© 2011 Michael Karpa

experience to understand a source text. This could be the greatest skill that translators have.

If understanding matters, that is. If it is possible to translate without understanding, the eventual reign of computers over all but the most celebrated works of art seems assured. So I sought to find out more about what it means to understand, what happens in the brain when we do, and if that could help explain why Japanese-to-English translation is one of the language combinations in which computer translation has been, so far, among the least successful.

How Does the Brain Translate?

The essence of translation is reproducing as complete an understanding as possible in a different language. What does that entail in terms of brain function?

Let's say I am going to translate a sentence. First, I look at my text, a simple workaday source text: データベースの検索手順を次の図に示します。I say to myself, it's Japanese. Then I read it. Well, actually I scan it first to get a feel: データベース, 手順. It's a familiar subject, instructions about software. I can read it fluently and just take in the meaning. It refers to something that is next, a figure, so I look at the figure. I see how long the procedure is, who does what. It is time for a test translation. I disengage from Japanese and switch to thinking in English. *Database*. No, *procedure* should come before *database*. No, *search procedure*. A database search procedure, or a procedure for searching a database. And a *next figure*, which is *shown there*. So that's what it says: *The database search procedure is shown in the next figure*. My translator experience rejects this as too Japanese, too faux-English. *The next figure shows the database search procedure*. Or is it *the following figure*, or *the figure below*? Is it *shown* or *illustrated* or *depicted*? I decide: *The following figure shows the procedure for searching a database*. Is it *the database*? I have to read a bit more for that one, both before and after my text. A database it is.

データベースの検索手順を次の図に示します

The figure below shows the procedure for searching a database.

That's a lot going on for one short, simple sentence, and yet I do this all the time, nearly automatically, usually entirely unaware of the process. I identify the language, switch into that language, arrive at an understanding, refine that understanding, and test my understanding by checking contextual information. Once I feel I have an understanding that is accurate enough to build on, I switch languages to start testing translations. I try on subjects and objects, transitive and intransitive verbs, passive and active voices, refining and again checking the language for contextual appropriateness and register. I put down my favorite version and move on, knowing I'll look at it again later purely as English, without engaging the Japanese at all.

Where in the brain does all this occur? Is it in working memory, that part of the brain so overloaded by the Web? Theodore Zanto describes working memory as a cognitive action that allows us to keep information temporarily in mind after its source is gone, and to manipulate it as well, which is the widely accepted model proposed by Alan Baddeley.³ When I translate, I am certainly doing this. Reading the Japanese loads the sentence into my working memory, where I keep it refreshed by deliberate attention, a process called "rehearsing." I do this using a part of working memory called the central executive, which establishes and maintains goals and has been identified with the dorsolateral prefrontal cortex. In fact, I have been engaged far more deeply than I would be simply deciding whether to click a link. The central executive of my brain has been performing repeated executive tasks—switching languages on, switching them off, ordering words to be produced, associating information, and generating thoughts. Assuming the Internet hasn't hopelessly clogged up my dorsolateral prefrontal cortex, I appear to have been engaging in creative thought.

How Do We Read?

In translating, that process begins with reading, which in turn begins with recognizing graphic symbols as meaningful. Reading is not innate to the human brain. It is a skill that must be painstakingly taught and learned. Reading stories from picture books amazes young children with adults' ability to

reproduce a story, time after time. Repetitive readings of rhythmical, rhyming verses also stealthily introduce children's growing minds to the idea that words are composed of phonemes—discrete auditory units. This lays the foundation for reading: the recoding of graphically recorded phonemes, regardless of orthography, into speech and language.

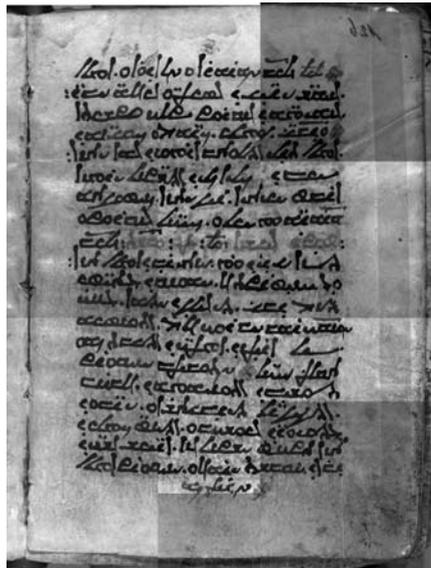
But writing did not begin with phonemes. The first event in the dawn of writing was symbolic representation (numbers on clay tokens), the second was the insight that symbols can represent words, and the third was that symbols could correspond to sounds.⁴ The earliest evolving orthographies—hieroglyphics, cuneiform, Chinese, and Zapotec—were logographic, representing words using graphic symbols. These graphs were originally independent of sound, but soon came to do double duty, representing words in some cases (logographs) and sounds in others (syllabaries), just as the kanji precursors of kana did (奈 evolving into な). These writing systems were logosyllabaries, composed of both logographs and phonetic symbols, as well as characters with features of both (think 清 and 請).

Early logographic and logosyllabic writing systems were extremely time-consuming to learn and, therefore, the domain of only an educated few. Such systems had far fewer characters than today, so characters served multiple roles, which required notations to mark the particular role a character was serving. The complexity was reminiscent of Japanese written in 漢文, with its 訓点, 返り点, 句読点, and so on. Alphabets were a democratizing invention, but their initial absence of vowels still meant considerable skill was required to read. Even after the Greek invention of vowels, reading was arduous because there were no spaces between words, a style of writing called *scriptura continua*, which Paul Saenger describes in great detail in *Space Between Words: The Origins of Silent Reading*.⁶

A key technique, one readers still use today, made this possible: reading aloud. The difficulty in reading without spaces is that words are tough to identify. (This is also, crucially, a reason given for the failure of machine translation and translation memory programs to optimally translate Japanese or Chinese into alphabetic languages: the lack of spaces in Japanese and Chinese makes

words hard to identify.) So how did humans read the unseparated stream of *scriptura continua*? They had to scan back and forth, reading syllables aloud to assemble an understanding. Physically pronouncing ambiguous phonemes helped readers keep them in mind until the larger words and sentences became clear. This is true in early learning of modern Japanese as well. According to Takahiko Sakamoto and Kiyoshi Makita, kana imposes a heavy burden on young readers by making them reconstruct words from phonetic signs that offer no distinction between words boundaries. “To aid the young Japanese reader, group oral recitation is essential.”⁷

There is a reason why reading aloud works: one of the four parts of working memory is based in sound. Alan Baddeley’s model of working memory is that of three slave components supervised by a central executive: a phonological loop in the left (language-dominant) hemisphere, a visuo-spatial sketchpad in the right hemisphere, and an episodic buffer. The phonological loop holds auditory memories. These decay rapidly unless rehearsed (repeatedly articulated within working memory). The visuo-spatial sketchpad stores visual information (form and color), deals with spatial and movement information, and likewise rehearses information. The episodic buffer helps integrate visual, spatial, and verbal information with time sequences, such as memories of stories, and is linked to long-term memory and semantic meaning. This information is then transferred to the central executive.⁸ Back in the days when translators used multiple, heavy dictionaries stored on shelves, many of us used a trick to cut down on reaches to the shelf: when looking up two or more Japanese terms, say the first English translation aloud to keep it in memory while look-



Ninth-century translation into Syriac of a Greek medical text by Galen. Written in *scriptura continua*.

ing up the second (which you then remember visually). Without knowing it, we were taking advantage of the phonological loop to rehearse the first term while loading the second into the visuo-spatial sketchpad.

In his landmark *Reading in the Brain*, Stanislaus Dehaene gives an elegant summary of how the brain turns written symbols into language in the reading process.

Written word processing starts in our eyes. Only the center of the retina, called the fovea, has a fine enough resolution to allow for the recognition of small print. Our gaze must therefore move around the page constantly. Whenever our eyes stop, we only recognize one or two words. Each of them is then split up into myriad fragments by retinal neurons and must be put back together before it can be recognized. Our visual system progressively extracts graphemes, syllables, prefixes, suffixes, and word roots. Two major parallel processing routes eventually come into play: the phonological route, which converts letters into speech sounds, and the lexical route, which gives access to a mental dictionary of word meanings.⁹

With unseparated text, it was impossible to do this in silence. The eyes also had more to do. In reading, they move from stop to stop (fixations) in a series of jumps (called “saccades”). Without spaces to mark out the words, ancient readers needed more than double the number of fixations and saccades we do today. Ancient readers “also required a quantity of ocular regressions for which there is no parallel under modern reading conditions in order to verify that the words had been correctly separated.”¹⁰

Reading was inefficient and laborious. It was work. The notion of reading for pleasure or for reference was unknown. More typically, texts were “performed” by professionals. The difficulty of reading *scriptura continua* resulted in the delegation of reading and writing to skilled slaves.

The medieval introduction of spaces into Latin and Greek *scriptura continua* (first in Ireland, then England, and then the European continent) created the possibility of fluent reading, since a major part of the work was already

done for the reader. Word order was also becoming more important. Highly inflected classical Latin did not place importance on word order, except for metrical and rhythmical eloquence, much as the use of particles in Japanese allows phrases marked by particles to be moved with some freedom within Japanese sentences. At the end of the Imperial Roman period, however, word endings were increasingly not being pronounced, which obscured inflections and made sequence more important in Vulgar Latin (unwritten or colloquial Latin). Word order began to follow fixed sequences. This helped make silent reading possible, which made reading easier, and thus more efficient.¹¹

In Japanese and Chinese, a “lack of spaces” clearly does not impede reading. Saenger remarks that, indeed “many skilled adult Chinese readers are able to achieve a proficiency in rapid, silent reading perhaps unequaled in modern Occidental languages.” Dehaene notes that the greater density of characters in Chinese results in shorter saccades (jumps of focus).¹² Individual characters function much more like words than letters, so the eye does not need to jump around as much. The way kanji and kana are combined helps readers of Japanese pick out separate words. A transition to kanji from kana frequently marks the beginning of a word, so the reader is not forced to sound out phonetic elements and identify each and every word through a hypothesis-test procedure.

The result in both cases is silent, automatic reading, which is mentally much more efficient. Now that I as the translator have phonologically recoded the orthography efficiently into language (that is, I have read the Japanese sentence), I rehearse it in working memory. Next comes the fun part: I switch languages and reproduce my text in English.

The Bilingual Brain

So how do brains work with multiple languages? There is no simple answer. Bilingual brains may have two different types of structures. One functional magnetic resonance imaging (fMRI) study from 1997¹³ suggested that early bilinguals (in this study, those exposed to two languages from infancy) produce both their languages in the same part of the brain, but that late bilinguals may use distinct regions for the dominant and second languages. In

late bilinguals, semantic processes may play a greater role than phonological processes.

A 1999 study by Cathy Price, David Green, and Roswitha von Studnitz¹⁴ of language switching and translation (of single words only) explained that the bilingual brain has a “bilingual lexico-semantic system” that processes the look (orthography), sound (phonology), and meaning (semantics) of words in the dominant language and the second language. Price and colleagues made PET scans of very proficient German-English adult bilinguals while they were translating or reading words in German, English, or alternating languages. (Their German subjects learned English as their first foreign language at a mean age of 8.8 years and were fluent in English for a mean 9.8 years at the time of the experiment.)

In their work, they used the Inhibitory Control (IC) model of Green,¹⁵ which holds that language task schemas exist for each language. (A schema is an organized pattern of thought or behavior stored in long-term memory.) The schemas are external to the language system. A schema in the dominant language (say, a model for word production) will compete dynamically with a similar schema for the second language. To say something in the second language, you must inhibit the schema for the first. When you switch between languages, you inhibit the active schema and activate the other (previously inhibited) schema.

Price and colleagues found that when people switched languages, the areas of the brain that were activated were associated with phonological recoding, not semantic processing (turning written language into sound, rather than meaning).

I have noted in my own translation practice that translation sometimes flows nearly effortlessly. Other times, complex grammar, ambiguity, or unfamiliar subject matter slow me down. For the tougher sentences, I may return from intense concentration on my unfinished test translation in English only to find the Japanese original newly impenetrable. The harder I work creating a translation hypothesis (and the later in the day), the harder it becomes to switch. To get back in the Japanese groove, I may subvocalize or read aloud. The Japanese snaps into focus. It seems I have given phonological recoding a nudge.

Price and colleagues noted that psycholinguistic data emphasize two different routes for word translation: a direct route that links words at the syntactic level and an indirect route that links them by meaning. Their study found that translation increased activity in some regions involved in the control of action, though surprisingly not in the dorsal prefrontal cortex, possibly because of the high proficiency of the subjects. Perhaps with experience, single word translation becomes automatic. Nicholas Carr notes in his well-researched exploration of how the Web may impair deep reading, *The Shallows: What the Internet Is Doing to Our Brains*, that scans of experienced users of search engines show that they make use of the dorsolateral prefrontal cortex during search activity whereas inexperienced users do not.¹⁶ So perhaps experienced translators can leave their dorsal prefrontal cortexes free to be clogged by links.

In short, the Price et al. study confirmed that switching and translation use different parts of the language system. Translation uses semantics and articulation, language-switching uses phonology. The authors proposed that during translation the demands on articulation increase because the response to the input orthography must be inhibited while the response associated with the translation equivalent is activated.

This helps explain several issues with “word seeding,” the practice of globally replacing complex technical terms in a source document with their target-language translations. This step can help ensure terminological consistency, a sort of do-it-yourself translation memory program. I do this on occasion. It does free one from having to look up the same obscure term multiple times, or having to return to previous sections to see how you translated it before. And it prevents confusion when there are two very similar but crucially different terms. However, in practice I have found it sometimes counterproductive, because the seeded words do not seem to “function” with the words still in the source language. A perfectly good English adjectival phrase, comfortably settled in the form it will take in the final translation, nevertheless fails to modify its still-Japanese noun because the Japanese noun demands the features of a Japanese adjectival phrase. The two languages are in different schemas, which don’t always mingle easily. Or rather, if they are going to mingle, both must be

obeyed, much as a bilingual might say in a conversation with other bilinguals, “it was *きれいだった*.” Word seeding may also feel inefficient because of the switching costs of going from language to language. With word seeding, I may also lose confidence that the final translation is what was really meant because the original source term disappeared early on.

This problem turns up when translating with translation memory programs, as well. Statistical machine translation and translation memory programs have no real schemas. For these programs, language is not actively functioning. It obeys no rules. Rule-based machine translation, which codifies grammar, has paradoxically proven unsuccessful in actual translation. It lacks the huge store of experience a human mind contains, and semantics consequently elude it.

The strong role of semantics and articulation in translation may explain how we translate obscure or technical kanji terms we have never heard spoken and can't confidently pronounce—unexpected mixes of *訓読み* and *音読み* abound. In the absence of phonology, semantics carries the day. I have often felt there was a semantic route for kanji, a sort of *英読み* to take me directly from the logograph to the target language.

Philippa Jane Benson¹⁷ examined a landmark 1979 Stroop¹⁸ study by Biederman and Tsao that highlighted this idea. Biederman and Tsao compared Chinese and English Stroop interference and found significantly greater Stroop interference in Chinese subjects than in native English speakers. They speculated that Chinese readers may automatically process logographs by their configurations whereas readers of English use abstract rules to turn graphemes into sounds. They cited the widespread belief that Chinese characters provide more direct access to meaning than English words.¹⁹ But more recent research strongly indicates that reading of Chinese characters always involves phonological mechanisms, whether the reader is aware of them or not. In her study of Japanese acquisition of kana and kanji, Virginia A. Mann noted that, “phonetic representation is employed in the service of temporary memory and comprehension whether subjects are reading the English alphabet . . . or the Chinese logography.”²⁰

When I watch myself carefully, I realize that often as not, when I am reading a Japanese kanji phrase that I have never heard in use, I am either using a straightforward 音読み or, when desperate, pronouncing it in Chinese. (I studied Chinese for several years before starting Japanese, but have not kept it up.)

Dehaene notes that reading of both alphabetic and logographic writing activates the same area of the brain, which he calls the letterbox. Modern brain imaging has clearly established that, “in spite of geographical distance, different imaging methods, brain morphologies, education strategies, and writing systems, brain activations related to word recognition in Chinese readers lie only a few millimeters away from those of English readers.” Dehaene also notes that strokes that impair the letterbox affect reading of both kana and kanji in Japanese readers. However, there is more to language than the letterbox. Words written in kana and kanji do not activate the brain in identical ways. In rare cases, a stroke patient may retain the ability to read in one script but not in the other.²¹ The areas of the brain used for phonological processing also differ considerably between Chinese and English,²² and reading of characters makes extensive use of the right hemisphere, which alphabetic languages largely do not.

The key areas of difference for reading characters and alphabets are motoric memory and right hemisphere cortical regions. In addition to the semantic and phonological areas of the left hemisphere, reading of characters also recruits the visual-spatial system on the right hemisphere, which Li Hai Tan and colleagues have attributed to the need to analyze the square field of the character and determine the spatial locations of the brush strokes that comprise it.²³ Motoric memory is involved because Chinese and Japanese (and Sumerian) children learn (or learned) hundreds or thousands of logographic characters through years of repetitive writing in established stroke order. Mann describes her experience researching memory of graphic designs: “During testing, I noted that, unlike American children, many Japanese children attempted to trace the designs on their fingers or even with a motion of their head.”²⁴

Maryanne Wolf includes a figure illustrating the different areas of the brain activated while reading English, Chinese, and kana and notes that, unlike alphabetic writing systems, logographic characters require the involvement

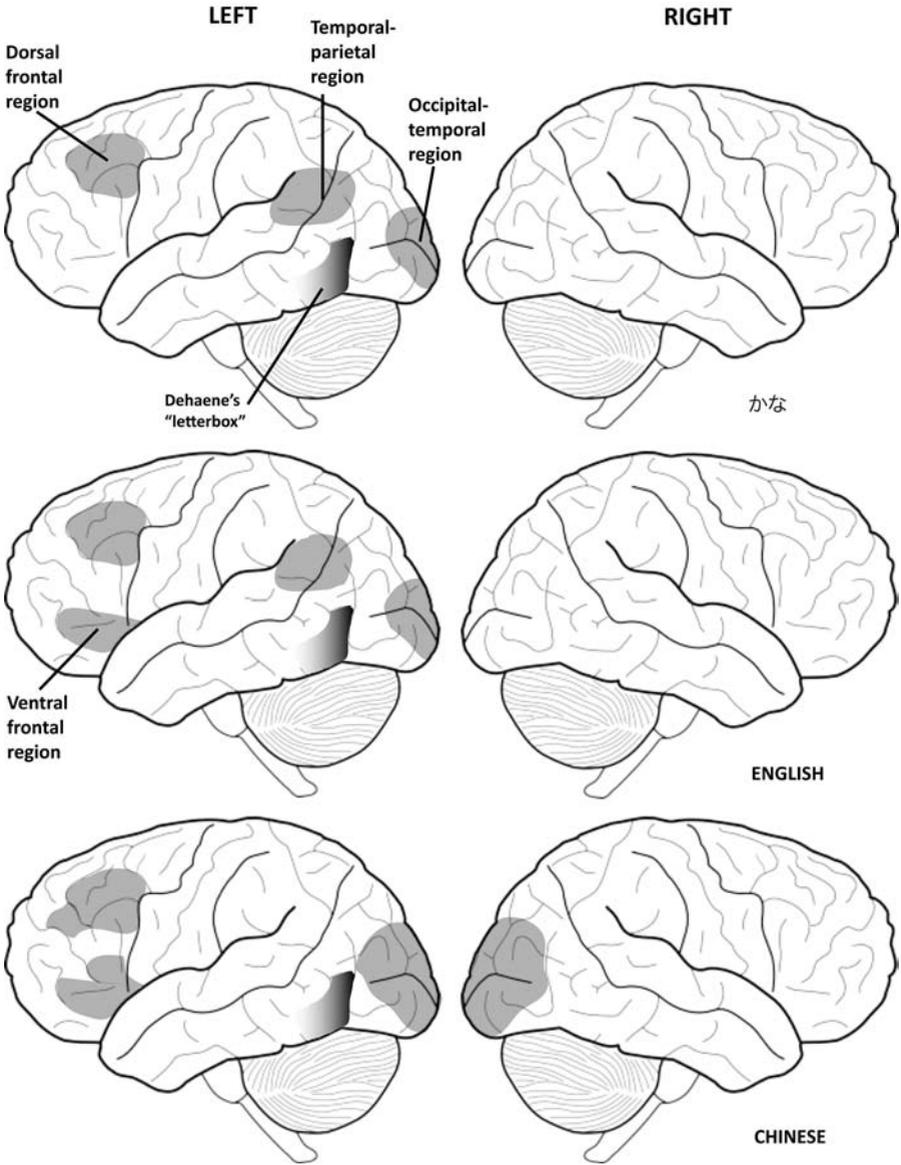
of right hemisphere areas known to contribute both to spatial analysis and to more global types of processing.²⁵

Understanding Understanding: The Conscious Translator

In translating Japanese to English we make use of semantics (meaning) and articulation (our ability to produce language), as well as motoric, visuo-spatial, and phonological memory. One thing that sets the human species apart is our frontal lobe. Humans have much more connectivity in our prefrontal neurons than other primates, we receive more incoming inputs, and we have many more synaptic contacts. We have one form of neuron, the giant fusiform cell, than we share only with other great apes. The extremely long axons of these neurons connect the frontal lobe to other distant regions of the cortex. Dehaene has proposed that this creates a “global neuronal workspace.” The prefrontal cortex is less specialized than most of the cortex but is connected by these giant fusiform cells to specialized areas throughout the brain. According to Dehaene, the main function of this “neuronal workspace” is:

to assemble, confront, recombine, and synthesize knowledge, ... guided by any combination of information from past or present experience. Our prefrontal cortex, thanks to its connections to all the high-level areas, provides a space for internal deliberations fed by a whole set of perceptions and memories. Bluntly put, what we term “conscious thinking” may simply be the manipulation of information within this global neuronal workspace.²⁶

Is this something that computers cannot reproduce? After all, our biological capabilities remain phenomenal: the total storage capacity of all the world’s computers is about the same as the DNA of one human, and the maximum number of instructions per second that all the world’s general computers could execute (in 2007) was equivalent to the maximum number of nerve impulses that one human brain can execute per second.²⁷ And computers do not have motoric memory, visuo-spatial, or phonological systems. At least, not yet. Odds are good that they will someday be able to reproduce a global



Reading in three languages: The brain uses different areas when reading a syllabary, alphabet or logographs. The approximate location of Dehaene's "letterbox" (visual word form area) is also shown. Diagram by Michael Karpa, based on Maryanne Wolf's diagram in *Proust and the Squid*, p. 62.

neuronal workplace. Rule-based MT has not been as successful as statistical MT, but it may yet find its purpose as a kind of language-task schema married to statistical MT, which will serve the role of long-term memory under a yet-to-be-invented central executive. If nothing else, science is showing us that—marvelous, complex, and flexible though it is—the brain is understandable. In an interview available on YouTube, Alan Baddeley says:

Consciousness is a way in which evolution has equipped us with a tool for pulling together information and allowing us to reflect on it and use it in order to plan for the future. . . . The episodic buffer . . . allows information to be combined, bound together from many sources . . . within a system that can stand above it and reflect on it and perhaps even reflect on the reflection. So, I think of consciousness as a biological function, a function that we're starting to understand.²⁸

In a sense, if we ask what is the value of a human translator, we ask what is the value of consciousness. Consciousness, the ability to reflect on our thinking, should someday be creatable. What our ability to understand really offers is the ability to choose what we do with our consciousness. It is up to us as individuals how good a job we do as translators, or as anything.

The advance that resulted from introducing spaces into *scriptura continua* was not speed, but efficiency. It freed our minds to think as we read. The greater efficiency of automatic, silent reading allows associative memory to function even as we read. To the recorded words we add our own associations, memories, and emotions, generating new thoughts of our own.

Making the task of translation easier could also create mental room for new associations, new thoughts, and greater creativity, allowing us to produce better work. As a translator today, I wonder, could the rigid norms of translation memory programs actually lead to an increase in neurophysiological efficiency similar to that which occurred with the switch to silent reading?

At the moment, I sense not. Or at least, not yet. MT cannot translate on its own; it can only skillfully retrieve the work done by translators in the past.

Franz Och, the head of Google's machine translation team, has said Google's "algorithms basically mine everything that's out there."²⁹ With so little regard displayed for intellectual property rights and, by extension, the value of an individual translator's work, one might be forgiven for feeling that the skilled slave is an alternative model of the future of the translation profession. But uncertainty remains innate in language. We cannot predict what will happen between the beginning and the end of a short sentence, so translators have to be conscious. Our ability to stay conscious and present and to bring our wealth of experience to bear is what keeps us from becoming slaves. In the future, we may have another way to view a real computer translator, when it finally arrives: not as a tool or a competitor, but as something evolving into a collaborator, whose consciousness can work in tandem with our own profuse and unpredictable ability to think.

¹ "Translating in the Deep End" discussed how Google Translate and the cognitive effects detailed by Nicholas Carr in *The Shallows: What the Internet Is Doing to Our Brains* (W.W. Norton, New York, 2010) are affecting translators and the translation industry (*The ATA Chronicle*, American Translators Association, Alexandria VA, Jan 2011).

² This proposition was given wide exposure in Maryanne Wolf's 2007 work, *Proust and the Squid: The Story and Science of the Reading Brain* (New York: Harper). It continues to be documented by many others, including Elias Aboujaoude in *Virtually You: The Dangerous Powers of the E-Personality* (New York: W.W. Norton, 2011) and the laboratory of Adam Gazzaley at the University of California San Francisco.

³ Theodore P. Zanto, Michael T. Rubens, Arul Thangavel, and Adam Gazzaley, "Causal Role of the Prefrontal Cortex in Top-down Modulation of Visual Processing and Working Memory," *Nature Neuroscience* (advance online publication), 27 March 2011.

⁴ Wolf, *Proust and the Squid*, pp. 25–26.

⁵ About 85 percent of Chinese characters today contain a phonetic component to hint at pronunciation, but only 28 percent of these phonetic components actually sound the same as the whole characters that contain them. (Li Hai Tan, Angela R. Laird, Karl Li, and Peter T. Fox, in "Neuroanatomical Correlates of Phonological Processing of Chinese Characters and Alphabetic Words: A Meta-Analysis," *Human Brain Mapping* 25 (2005), pp. 83–91.

⁶ Paul Saenger, *Space Between Words: The Origins of Silent Reading* (Stanford University Press, 1997).

⁷ Takahiko Sakamoto and Kiyoshi Makita, in John Dowling, ed., *Comparative Reading: Cross-national Studies of Behavior in Experimental Psychology* (New York: Macmillan, 1973), pp. 440–65, as quoted by Saenger in *Space Between Words*.

⁸ Alan D. Baddeley and Robert H. Logie, "Working Memory: The Multiple Component Model," in Akira Miyake and Priti Shah, eds. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (Cambridge University Press, 1999), pp. 28–61.

⁹ Stanislas Dehaene, *Reading in the Brain: The Science of Evolution of a Human Invention* (New York: Viking Press, 2009), p. 11.

- ¹⁰ Saenger, p. 7. There is perhaps a parallel in the world of Japanese-English translation, in which we go through a process of hypothesis-test of potential translations of long Japanese sentences, scanning back in the Japanese as we go to test whether our translation is working out.
- ¹¹ Around the year 1750, there was another dramatic change in the way people read documents. Printing greatly increased the volume and breadth of available books. Before this time, people read intensively: they read the same few books over and over. By the early 1800s, however, people started to read things extensively. See Ziming Liu, "Reading Behavior in the Digital Environment: Changes in Reading Behavior over the Past 10 Years," *Journal of Documentation* 61:6 (2005), pp. 700–12.
- ¹² Dehaene, *Reading in the Brain*, p. 17.
- ¹³ Karl H. S. Kim, Norman R. Relkin, Kyoung-Min Lee, and Joy Hirsch, "Distinct Cortical Areas Associated with Native and Second Languages," *Nature* 388 (July 1997), pp. 171–44. A later fMRI study showed a pattern of overlapping activations that was similar in early and late bilinguals (Michael W. L. Chee, Edsel W. L. Tan, Thorsten Thiel, "Mandarin and English Single Word Processing Studied with Functional Magnetic Resonance Imaging," *Journal of Neuroscience* 19 (1999), pp. 3,050–56).
- ¹⁴ Cathy J. Price, David W. Green, and Roswitha von Studnitz, "A Functional Imaging Study of Translation and Language Switching," *Brain: A Journal of Neurology* 122:12 (1999), pp. 2,221–35.
- ¹⁵ David W. Green, "Control, activation, and resource: a framework and a model for the control of speech in bilinguals," *Brain and Language* 27 (March 1986), pp. 210–23.
- ¹⁶ Carr, *The Shallows*, p. 121.
- ¹⁷ Philippa Jane Benson, "Cross-orthographic Stroop Research: One Study in Context," *Sino-Platonic Papers* 21, ed. Victor H. Mayor, University of Pennsylvania, December 1990.
- ¹⁸ Stroop interference is the time difference between the presentation of a color-word and the subject's response to a color naming or word-reading instruction when a color word is printed in ink of a different color (e.g., the word *green* printed in red ink).
- ¹⁹ According to Price and colleagues, the dominant language is more strongly inhibited, so it requires more time to be reactivated.
- ²⁰ Virginia A. Mann, "Temporary Memory for Linguistic and Non-linguistic Material in Relation to the Acquisition of Japanese Kana and Kanji," in Henry S.R. Kao and Rumjahn Hoosain, eds., *Linguistics, Psychology, and the Chinese Language* (Hong Kong: University of Hong Kong Press, 1986), pp. 55–167.
- ²¹ Dehaene, *Reading in the Brain*, pp. 97–100.
- ²² Li Hai Tan, et al., in "Neuroanatomical Correlates of Phonological Processing of Chinese Characters and Alphabetic Words: A Meta-Analysis".
- ²³ Li Hai Tan, Ho-Ling Liu, Charles A. Perfetti, John A. Spinks, Peter T. Fox, and Jia-Hong Gao, "The Neural System Underlying Chinese Logograph Reading," *Neuroimage* 13 (2001), pp. 836–46.
- ²⁴ Mann, "Temporary Memory," pp. 131–32.
- ²⁵ Wolf, *Proust and the Squid*, pp. 35–36. See also the discussion of use of the two hemispheres in reading single characters and two-character words in Li Hai Tan, John A. Spinks, Jia-Hong Gao, Ho-Ling Liu, Charles A Perfetti, Jinhui Xiong, Kathryn A. Stofer, Yonglin Pu, Yijun Liu, and Peter T. Fox, "Brain Activation in the Processing of Chinese Characters and Words: A Functional MRI Study," *Human Brain Mapping* 10 (2000), pp. 16–27.
- ²⁶ Dehaene, *Reading in the Brain*, p. 318.
- ²⁷ Martin Hilbert and Priscila López, "The World's Technological Capacity to Store, Communicate, and Compute Information," *Science* (Feb. 11, 2011), pp. 692–93.
- ²⁸ <http://www.youtube.com/watch?v=pAcqWd4X50g>
- ²⁹ Interview with Franz Josef Och in "Technology," David Sarno, March 11, 2010, *Los Angeles Times*.