

Interview with Mike Dillinger, Machine Translation Pioneer

Translation has been an integral part of our human evolution ever since we developed the ability to speak. But with the arrival of the industrial age, the spread of globalization, and the birth of technology, it was inevitable that someone would eventually wonder if machines were capable of performing this task. Research began in earnest in the early 1950s, and great progress has obviously been made since then. Machine translation, or MT, is the name for both the technology and for an established subfield of linguistic engineering that creates systems to translate text or speech from one human language to another.

My guest this time has spent some 20 years deeply immersed in the theory and practice of MT. Mike Dillinger currently manages taxonomies and MT at LinkedIn. He is a member of the External Advisory Board at the ADAPT Centre (a leading center for MT research) in Dublin, Ireland, and advises startups in the U.S., Israel, Australia, and Brazil. He has twice been president of the Association for Machine Translation in the Americas, and has experience with all phases of planning, development, deployment, and evaluation of MT systems. He led LinkedIn's and eBay's first launches of full-scale production MT, spearheaded the development of the first commercial MT-translation memory (TM) tool integration (Star Transit with Logos MT), and developed an interactive speech-to-speech MT system.

Mike earned a PhD from McGill University, in Montreal, for research on the cognitive processes of simultaneous interpreting and comprehension of technical content, as well as degrees in linguistics in the U.S., Canada, and Brazil. He is also an experienced translator and interpreter who has worked in English, Portuguese, Spanish, and French. Mike wants organizations everywhere to enable global communication by developing content that's understandable and translatable, and by deploying MT effectively.



Mike Dillinger

Machine translation, or MT, is the name for both the technology and for an established subfield of linguistic engineering that creates systems to translate text or speech from one human language to another.

Thank you for spending time with us, Mike. Let's start with a little history. Can you summarize the highlights of MT's development?

A research collaboration between Georgetown University in Washington, DC, and IBM resulted in the first widely-known MT system from Russian into English. It was reported on the front page of *The New York Times* on January 8, 1954, and, of course, led to the hasty conclusion that the problem of automatic translation was essentially solved.¹

Fun facts: it wasn't a real MT system, just a proof of concept, with about two dozen rules and fewer than 200 words in its dictionary. It was basically what we now call a hybrid system, in which rules used the most probable word senses. It was, however, a spectacularly

effective system because it convinced the U.S. and Russian governments to pour millions of dollars into research to "finish" developing the technology. The U.S. first used MT to translate Russian scientific publications to track their technology development during the Cold War. Later, MT (from Logos Corp) was used to translate helicopter repair manuals into Vietnamese.

Fast forward through the 1970s, when U.S. government funding dried up and Europe and Japan took the lead in MT research, trying to capture translators' knowledge in rules stored on a computer—"rule-based MT." In the 1980s, again at IBM, a new statistical approach emerged in which the software would calculate the likelihood of one particular translation based on many examples of human translations—"statistical MT." Again, this new approach convinced the U.S. government to provide financial support, and MT took off as a research area in the 1990s.

In the early 2000s, Google formed a group to work on MT, and companies like Language Weaver started to offer statistical MT systems as products. Since 2007, much progress has been made in making MT technology more easily available, both to researchers (with the Moses toolkit) and to translators, with Microsoft Translators Hub, Google's Translator Toolkit, and products like Kantan MT. Most recently, significant progress has been made on two fronts:

1. MT systems can now be "adaptive." In other words, the systems get updated every time a human makes a correction. (One example of this can be found at www.lilt.com.)
2. A newer research approach called "neural" MT is improving how MT systems leverage information in the context of a sentence being translated.

Most translators are, of course, very familiar with MT, and use it in their daily work in one way or another. Just to get us all on the same page, please give us your definition of MT: a description of how it works and how it differs from Translation Memory (TM).

Both technologies (TM and MT) have the same job: to provide possible translations for expert review. TM technology focuses on reusing whole segments (mostly sentences). But you don't get anything if most of an incoming sentence doesn't match, unless you ask for an "assembled" translation that guesses piece by piece. MT just produces assembled translations, but in a much more sophisticated way. That's it.

Speed is obviously one of MT's great advantages. How fast can a machine translate, compared to an average human translator?

The average human translator delivers something like 2,000 words *per day*. An average MT system can produce

something like 2,000,000 words per day, and engineers know how to increase that rate with super-fast TMs.

Can you give us an idea of the volume involved? How many words are being processed by MT programs on an average day in the U.S. and worldwide? Is it possible to compare that with the volume of human translation?

Let's take one example: Google Translate. It processes about 100 billion words per day, in 103 languages, for 500 million users around the world. That's equivalent to the output of 50 million translators every day.

Translation technologies are translator "accelerators," not translator "replacements."

Where is MT mainly used? In what areas, what sectors?

MT is mainly used in two scenarios:

1. The "no-other-choice-but-MT" scenarios (e.g., translating search queries, e-mail, tweets, and user feedback), ecommerce, informal "I'm-just-curious translations," translations of customer support information, and espionage. What they all have in common is that the volume of source documents is far too large for humans to tackle, and the information in the source documents is too ephemeral or not very valuable. It's also really horrible translation work. Can you imagine translating search queries or e-mail all day! In this sense, MT is doing us a favor! This kind of MT actually creates *more* work for human translators. It helps organizations identify important information—serving a triage function—that humans often have to translate.
2. The "I'm-in-a-hurry" scenarios, such as big localization projects for multinational contracts and project launches at global companies. Again, there's usually far too much content and too little time for normal

translation processes, so we use MT to pre-translate. When we set things up correctly, MT can make the whole process about four times faster.

How has MT impacted technical writing, and source-material writing in general?

I've worked quite a bit with tech writers at a range of companies, and it's clear to me that the impact of MT itself has been very small. The impact of TM technology and translation pricing, however, has been substantial. To save money (reason #1!), decrease project turnaround time, and increase readability for end users, many companies have adopted tools like Acrolinx and content management systems to improve the consistency and reuse of their source content.

In the early days of TM and MT, some translators were apprehensive, scared that technology was about to make them redundant. You had a great response to that. You said: translation technologies are translator "accelerators," not translator "replacements." Please expand on that idea.

Translation technologies today perform pretty poorly for most kinds of content, especially when writers don't write consistently or clearly. TM and MT tools simply aren't mature enough to be "let out of the house" on their own when the goal is publishable content. In any scenario where the source content is valuable or the reader's understanding is important, we continue to need human translators. By the way, the amount of economically valuable content that should really be translated is estimated to be at least 10 times what we are doing today. That means that we need to find ways to accelerate the translator's productivity, so we use MT as a translator accelerator.

You have also said that there are two ways for translators to work with MT: post-editing and turbo-translating. Please tell us more about those two methods.

In post-editing, someone else controls the MT system that provides you with draft translations. Often, these individuals don't know what they're doing or aren't paying attention to the same things as you. So, you don't always get the best possible candidate

Simultaneous Interpretation Equipment

FM Transmitters Receivers

Sales & Rentals

Affordable Pricing Superior Sound

**5 YEAR WARRANTY
30 DAY MONEY BACK
GUARANTEE!**

**(888) 677-4387
(518) 426-0171**

LANDMARK AUDIO FM TECHNOLOGIES Systems

www.LandmarkFm.com

translations with which to work. That's why you have to be extra careful when you accept post-editing jobs, since not all MT output is of the same quality.

In "turbo-translating" you know enough about MT to control the process and create your own draft translations. This means you can better manage the resources the system uses, you control the trade-off between speed and accuracy, and you can predict the kinds of problems you'll find and prepare for them. The older rule-based MT systems offered more direct control over different linguistic parameters, so they were better for turbo-translating. Statistical systems are usually black-box systems that we can't control much.

Actually, I think there is a third way to work with MT, but it's not available yet. I call this new option "hybrid-intelligence" translation. The idea is to leverage the strengths of both machine intelligence and human expertise by letting the humans "drive" the machine. This approach is like fly-by-wire systems for pilots: the pilot is definitely in charge and the system works out thousands of routine details to allow the pilot to focus on the important things. Adaptive MT is the first step in this direction, and I think there are many more things we can do to allow translators to "pilot" their MT systems.

The post-editing field looks like a future growth area for many in our profession. Is that how you see it? Is post-editing something any translator can do? What are the requirements?

Yes, post-editing will surely continue to grow, and very rapidly. There's some controversy concerning your second question. I believe that anyone who can correct a junior translator's work can correct MT output. Researchers have already documented huge differences in how quickly people can do post-editing, and a large part of that is due to experience rather than special training. I still can't identify any specific training that post-editing requires.

Do all post-editing clients want the same "product" from a translator, or are there different standards or levels? Do they all want a translator's best work, or do some want something that's just "good enough"?

This is actually the beauty of post-editing. Whereas before MT we could either provide a first-rate translation or none at all, now we can calibrate the translation quality more precisely to the client's needs. For a while, we saw people ask for "light" post-editing and "full" post-editing. "Light" post-editing is an effort to find and fix only the most misleading and blatantly incorrect translations (e.g., missing negation). "Full" post-editing is the task of bringing MT output up to your usual, high standard for human translation quality. Unfortunately, it's hard for translators and clients to agree on when we're done with "light" post-editing, so it's a headache to manage. "Medium" post-editing has appeared as an option that's easier to manage: fix only terminology and grammar, and don't worry about style and tone versus "full" post-editing, where we have to fix everything.

We read that improvements in MT technology for spoken language applications are being driven by the interpreting requirements of military operations overseas. How will those improvements impact civilian interpreters working in our usual environments (medical, legal, corporate, etc.)?

Military operations overseas have incredibly demanding requirements for speech-to-speech MT. This technology has to interpret between uneducated speakers of unusual dialects of unheard-of languages and more educated speakers of varying dialects of English using machines with no Internet connection in extreme weather conditions. The equipment has to be light enough to carry along with a 40-pound backpack and robust enough for a truck to drive over, with batteries that last for weeks. And speech recognition has to work in the middle of traffic and gunfire. Oh, and we have to build the system with next to no example sentences (data collection in a war zone isn't easy) and make sure that it can cover a wide range of topics. If we can make progress on any of these fronts, then civilian interpreting technology will certainly improve. Right now, soldiers in the field don't have any MT systems to use. They rely on human interpreters.

Is there a healthy level of international cooperation in MT development? At what level does that occur: government, military, industry, or academia?

Yes! The International Association for Machine Translation holds its MT Summit every other year to gather together the global MT community. Researchers, both in industry and academia, collaborate routinely across national boundaries. Multinational companies hire people from around the world. I've consulted for the European Community and worked on MT projects in at least five countries.

How does a machine learn—and keep learning—how to translate? How does it incorporate new words, phrases, and terminology into its repertoire? How does it expand its ability to process syntactical shifts and other linguistic features as language evolves?

Machines "learn" by ingesting, analyzing, and storing information about example human translations. Notice the scare quotes: machines only "know" what they've seen. An MT system has a huge database of all the words it has ever seen, all the translations for each word it has ever seen, and all the contexts in which both the word and its translations have occurred. It "learns" by adding more human examples to this database and by recalculating the most likely translations for each source sequence. It evolves by acquiring more information about some words and sentences than about others. We have to feed the system with more example translations continuously. Lots and lots of example translations, until we get good coverage of the words and sentence types that we need for a specific project. And we need linguists to do this kind of "feeding" work.

As the volume of available content requiring translation increases exponentially daily, it's clear that human translators cannot supply the need. Can current levels of MT keep up?

MT can keep up in terms of quantity, but not quality. The challenge is the increase in *valuable* content, which MT can't handle well and is already too much for humans to handle. Today, a great deal of valuable content simply goes untranslated.

ANNOUNCEMENT

Upcoming ATA Webinars for 2017

www.atanet.org/webinars

April 4, 2017

Translation Contracts: Beyond the Basics

Time: 12:00 noon U.S. Eastern Standard Time

Registration: ATA Member: \$45; Non-Member \$60

This webinar will discuss non-disclosure agreements (when and under what circumstances to sign); the benefits of having your own terms of service; and how to use ATA's model agreement and a contract checklist to draft your own terms of service. We will also examine some of the most common "problem clauses" in translation contracts and learn how to interpret them.

May 4, 2017

Transcreation: Translation with a Twist

Time: 12:00 noon U.S. Eastern Standard Time

Registration: ATA Member: \$45; Non-Member \$60

This webinar provides some background on how transcreation works, which skills are required, and how you go about transcreating copy. Some English into Dutch examples will be provided, but understanding Dutch is not required.

May 25, 2017

Why Can't I Raise My Rates?

Time: 12:00 noon Daylight Saving Time

Registration: ATA Member: \$45; Non-Member \$60

This webinar will provide an introduction to the economics of language services. Freelance translators and interpreters have some level of control over what they charge, but the market is the ultimate arbiter of pricing. We'll explore why this is so and what freelancers can (and can't) do to influence it. Along the way, we'll discuss concepts such as value, scarcity, marginal utility, supply, demand, price determination, competitive markets, commodities, and differentiation as they pertain to the language services market.

Online and On Demand!

Strengthen your skills and expertise with the education you need when you need it. Effective and affordable, more than 60 ATA webinars are available for streaming 24/7. Check out the complete library at www.atanet.org/webinars.

CAN'T ATTEND?

Register for the webinar now and a link to the recorded version will be sent to you after the live event!

OUR WORLD OF WORDS continued

We read that 14 languages are required to reach 90% of the world's most economically active populations, but most websites can only deliver content in about seven. What are those seven languages? Which languages will be next as capacity increases?

These are what companies call the Tier 1 languages. Although the list varies from company to company, it usually includes English, Simplified Chinese, Spanish, Russian, Japanese, German, and French. The next batch, the Tier 2 languages, varies more according to each company's international strategy, but usually includes Korean, Arabic, Italian, Indonesian, Dutch, Traditional Chinese, and the Scandinavian languages.

Finally, please tell us what developments in MT we can expect to see in the near-to-medium future, and where translators and interpreters fit in that evolving scenario.

In my opinion, the most interesting areas of MT research are domain adaptation, neural MT, and hybrid-intelligence systems.

Domain adaptation is the part of an MT system that tries to pick and choose which translations, of all the millions of translations the system has seen, will be most relevant for your specific project. There's some really fascinating work going on in Europe to make this adaptation faster and more accurate so that we get much better candidate translations.

Neural MT, the hottest, latest fashion in research circles, uses much more contextual information far more efficiently. It promises far better word sense disambiguation so that we get more accurate word choice in the candidate translations it proposes. It's also producing grammatically better candidate translations. Neural MT is already showing up in online MT systems and many more improvements are sure to come.

Hybrid-intelligence translation systems will someday let the translator "drive." The main assumption is that for the foreseeable future, MT won't be able to do publication-level translation of valuable information on its own. So, we must find ways to merge the things that MT systems do well with the things that only human translators can do well. The first systems of this type we call adaptive MT, which is built bottom-up for translators: when you (or your team) correct a translation, your corrections are applied

immediately to the remaining unreviewed sentences. An adaptive MT system "learns" much more relevant and more reliable information (for that particular project), learns it much faster, and presents it back to the translator much more quickly than ever before. In future systems, translators will not only correct a machine's output; they will also teach it linguistic rules, stylistic preferences, and project-specific idiosyncrasies.

Thank you, Mike, for this highly illuminating review of your fascinating field. ☺

NOTES

¹ Plumb, Robert K. "Russian Is Turned into English By a Fast Electronic Translator," *The New York Times* (January 8, 1954), 1, <http://bit.ly/NYT-electronic-translation>.

Also see: Hutchins, John. "The Georgetown-IBM Experiment Demonstrated in January 1954" (Association for Machine Translation in the Americas), <http://mt-archive.info/AMTA-2004-Hutchins.pdf>.



Tony Beckwith was born in Buenos Aires, Argentina, spent his formative years in Montevideo, Uruguay, then set off to see the world. He moved to Texas in 1980 and currently lives in Austin, Texas, where he works as a writer, translator, poet, and cartoonist. Contact: tony@tonybeckwith.com.

Sources for More Information on Machine Translation

ADAPT Centre

www.adaptcentre.ie

Association for Machine Translation in the Americas

<https://amtaweb.org>

International Association of Machine Translation

www.eamt.org/iamt.php

Lilt

www.lilt.com