The New York Times Magazine | https://nyti.ms/2hMtKOn

The Great A.I. Awakening

How Google used artificial intelligence to transform Google Translate, one of its more popular services — and how machine learning is poised to reinvent computing itself.

By GIDEON LEWIS-KRAUS DEC. 14, 2016

Prologue: You Are What You Have Read

Late one Friday night in early November, Jun Rekimoto, a distinguished professor of human-computer interaction at the University of Tokyo, was online preparing for a lecture when he began to notice some peculiar posts rolling in on social media. Apparently Google Translate, the company's popular machine-translation service, had suddenly and almost immeasurably improved. Rekimoto visited Translate himself and began to experiment with it. He was astonished. He had to go to sleep, but Translate refused to relax its grip on his imagination.

Rekimoto wrote up his initial findings in a blog post. First, he compared a few sentences from two published versions of "The Great Gatsby," Takashi Nozaki's 1957 translation and Haruki Murakami's more recent iteration, with what this new Google Translate was able to produce. Murakami's translation is written "in very polished Japanese," Rekimoto explained to me later via email, but the prose is distinctively "Murakami-style." By contrast, Google's translation — despite some "small unnaturalness" — reads to him as "more transparent."

The second half of Rekimoto's post examined the service in the other direction, from Japanese to English. He dashed off his own Japanese interpretation of the opening to Hemingway's "The Snows of Kilimanjaro," then ran that passage back

through Google into English. He published this version alongside Hemingway's original, and proceeded to invite his readers to guess which was the work of a machine.

NO. 1:

> Kilimanjaro is a snow-covered mountain 19,710 feet high, and is said to be the highest mountain in Africa. Its western summit is called the Masai "Ngaje Ngai," the House of God. Close to the western summit there is the dried and frozen carcass of a leopard. No one has explained what the leopard was seeking at that altitude.

NO. 2:

> Kilimanjaro is a mountain of 19,710 feet covered with snow and is said to be the highest mountain in Africa. The summit of the west is called "Ngaje Ngai" in Masai, the house of God. Near the top of the west there is a dry and frozen dead body of leopard. No one has ever explained what leopard wanted at that altitude.

Even to a native English speaker, the missing article on the leopard is the only real giveaway that No. 2 was the output of an automaton. Their closeness was a source of wonder to Rekimoto, who was well acquainted with the capabilities of the previous service. Only 24 hours earlier, Google would have translated the same Japanese passage as follows:

> Kilimanjaro is 19,710 feet of the mountain covered with snow, and it is said that the highest mountain in Africa. Top of the west, "Ngaje Ngai" in the Maasai language, has been referred to as the house of God. The top close to the west, there is a dry, frozen carcass of a leopard. Whether the leopard had what the demand at that altitude, there is no that nobody explained.

Rekimoto promoted his discovery to his hundred thousand or so followers on Twitter, and over the next few hours thousands of people broadcast their own experiments with the machine-translation service. Some were successful, others

meant mostly for comic effect. As dawn broke over Tokyo, Google Translate was the No. 1 trend on Japanese Twitter, just above some cult anime series and the long-awaited new single from a girl-idol supergroup. Everybody wondered: How had Google Translate become so uncannily artful?

Four days later, a couple of hundred journalists, entrepreneurs and advertisers from all over the world gathered in Google's London engineering office for a special announcement. Guests were greeted with Translate-branded fortune cookies. Their paper slips had a foreign phrase on one side — mine was in Norwegian — and on the other, an invitation to download the Translate app. Tables were set with trays of doughnuts and smoothies, each labeled with a placard that advertised its flavor in German (*zitrone*), Portuguese (*baunilha*) or Spanish (*manzana*). After a while, everyone was ushered into a plush, dark theater.

Sadiq Khan, the mayor of London, stood to make a few opening remarks. A friend, he began, had recently told him he reminded him of Google. "Why, because I know all the answers?" the mayor asked. "No," the friend replied, "because you're always trying to finish my sentences." The crowd tittered politely. Khan concluded by introducing Google's chief executive, Sundar Pichai, who took the stage.

Pichai was in London in part to inaugurate Google's new building there, the cornerstone of a new "knowledge quarter" under construction at King's Cross, and in part to unveil the completion of the initial phase of a company transformation he announced last year. The Google of the future, Pichai had said on several occasions, was going to be "A.I. first." What that meant in theory was complicated and had welcomed much speculation. What it meant in practice, with any luck, was that soon the company's products would no longer represent the fruits of traditional computer programming, exactly, but "machine learning."

A rarefied department within the company, Google Brain, was founded five years ago on this very principle: that artificial "neural networks" that acquaint themselves with the world via trial and error, as toddlers do, might in turn develop something like human flexibility. This notion is not new — a version of it dates to the earliest stages of modern computing, in the 1940s — but for much of its history most computer scientists saw it as vaguely disreputable, even mystical. Since 2011,

though, Google Brain has demonstrated that this approach to artificial intelligence could solve many problems that confounded decades of conventional efforts. Speech recognition didn't work very well until Brain undertook an effort to revamp it; the application of machine learning made its performance on Google's mobile platform, Android, almost as good as human transcription. The same was true of image recognition. Less than a year ago, Brain for the first time commenced with the gut renovation of an entire consumer product, and its momentous results were being celebrated tonight.

Translate made its debut in 2006 and since then has become one of Google's most reliable and popular assets; it serves more than 500 million monthly users in need of 140 billion words per day in a different language. It exists not only as its own stand-alone app but also as an integrated feature within Gmail, Chrome and many other Google offerings, where we take it as a push-button given — a frictionless, natural part of our digital commerce. It was only with the refugee crisis, Pichai explained from the lectern, that the company came to reckon with Translate's geopolitical importance: On the screen behind him appeared a graph whose steep curve indicated a recent fivefold increase in translations between Arabic and German. (It was also close to Pichai's own heart. He grew up in India, a land divided by dozens of languages.) The team had been steadily adding new languages and features, but gains in quality over the last four years had slowed considerably.

Until today. As of the previous weekend, Translate had been converted to an A.I.-based system for much of its traffic, not just in the United States but in Europe and Asia as well: The rollout included translations between English and Spanish, French, Portuguese, German, Chinese, Japanese, Korean and Turkish. The rest of Translate's hundred-odd languages were to come, with the aim of eight per month, by the end of next year. The new incarnation, to the pleasant surprise of Google's own engineers, had been completed in only nine months. The A.I. system had demonstrated overnight improvements roughly equal to the total gains the old one had accrued over its entire lifetime.

Pichai has an affection for the obscure literary reference; he told me a month earlier, in his office in Mountain View, Calif., that Translate in part exists because not everyone can be like the physicist Robert Oppenheimer, who learned Sanskrit to

read the Bhagavad Gita in the original. In London, the slide on the monitors behind him flicked to a Borges quote: "*Uno no es lo que es por lo que escribe, sino por lo que ha leído*."

Grinning, Pichai read aloud an awkward English version of the sentence that had been rendered by the old Translate system: "One is not what is for what he writes, but for what he has read."

To the right of that was a new A.I.-rendered version: "You are not what you write, but what you have read."

It was a fitting remark: The new Google Translate was run on the first machines that had, in a sense, ever learned to read anything at all.

Google's decision to reorganize itself around A.I. was the first major manifestation of what has become an industrywide machine-learning delirium. Over the past four years, six companies in particular — Google, Facebook, Apple, Amazon, Microsoft and the Chinese firm Baidu — have touched off an arms race for A.I. talent, particularly within universities. Corporate promises of resources and freedom have thinned out top academic departments. It has become widely known in Silicon Valley that Mark Zuckerberg, chief executive of Facebook, personally oversees, with phone calls and video-chat blandishments, his company's overtures to the most desirable graduate students. Starting salaries of seven figures are not unheard-of. Attendance at the field's most important academic conference has nearly quadrupled. What is at stake is not just one more piecemeal innovation but control over what very well could represent an entirely new computational platform: pervasive, ambient artificial intelligence.

The phrase "artificial intelligence" is invoked as if its meaning were self-evident, but it has always been a source of confusion and controversy. Imagine if you went back to the 1970s, stopped someone on the street, pulled out a smartphone and showed her Google Maps. Once you managed to convince her you weren't some oddly dressed wizard, and that what you withdrew from your pocket wasn't a black-arts amulet but merely a tiny computer more powerful than the one that guided Apollo missions, Google Maps would almost certainly seem to her a persuasive example of "artificial intelligence." In a very real sense, it is. It can do things any

map-literate human can manage, like get you from your hotel to the airport — though it can do so much more quickly and reliably. It can also do things that humans simply and obviously cannot: It can evaluate the traffic, plan the best route and reorient itself when you take the wrong exit.

Practically nobody today, however, would bestow upon Google Maps the honorific "A.I.," so sentimental and sparing are we in our use of the word "intelligence." Artificial intelligence, we believe, must be something that distinguishes HAL from whatever it is a loom or wheelbarrow can do. The minute we can automate a task, we downgrade the relevant skill involved to one of mere mechanism. Today Google Maps seems, in the pejorative sense of the term, robotic: It simply accepts an explicit demand (the need to get from one place to another) and tries to satisfy that demand as efficiently as possible. The goal posts for "artificial intelligence" are thus constantly receding.

When he has an opportunity to make careful distinctions, Pichai differentiates between the current applications of A.I. and the ultimate goal of "artificial *general* intelligence." Artificial general intelligence will not involve dutiful adherence to explicit instructions, but instead will demonstrate a facility with the implicit, the interpretive. It will be a general tool, designed for general purposes in a general context. Pichai believes his company's future depends on something like this. Imagine if you could tell Google Maps, "I'd like to go to the airport, but I need to stop off on the way to buy a present for my nephew." A more generally intelligent version of that service — a ubiquitous assistant, of the sort that Scarlett Johansson memorably disembodied three years ago in the Spike Jonze film "Her"— would know all sorts of things that, say, a close friend or an earnest intern might know: your nephew's age, and how much you ordinarily like to spend on gifts for children, and where to find an open store. But a truly intelligent Maps could also conceivably know all sorts of things a close friend wouldn't, like what has only recently come into fashion among preschoolers in your nephew's school — or more important, what its users actually want. If an intelligent machine were able to discern some intricate if murky regularity in data about what we have done in the past, it might be able to extrapolate about our subsequent desires, even if we don't entirely know them ourselves.

The new wave of A.I.-enhanced assistants — Apple's Siri, Facebook's M, Amazon's Echo — are all creatures of machine learning, built with similar intentions. The corporate dreams for machine learning, however, aren't exhausted by the goal of consumer clairvoyance. A medical-imaging subsidiary of Samsung announced this year that its new ultrasound devices could detect breast cancer. Management consultants are falling all over themselves to prep executives for the widening industrial applications of computers that program themselves. DeepMind, a 2014 Google acquisition, defeated the reigning human grandmaster of the ancient board game Go, despite predictions that such an achievement would take another 10 years.

In a famous 1950 essay, Alan Turing proposed a test for an artificial general intelligence: a computer that could, over the course of five minutes of text exchange, successfully deceive a real human interlocutor. Once a machine can translate fluently between two natural languages, the foundation has been laid for a machine that might one day "understand" human language well enough to engage in plausible conversation. Google Brain's members, who pushed and helped oversee the Translate project, believe that such a machine would be on its way to serving as a generally intelligent all-encompassing personal digital assistant.

What follows here is the story of how a team of Google researchers and engineers — at first one or two, then three or four, and finally more than a hundred — made considerable progress in that direction. It's an uncommon story in many ways, not least of all because it defies many of the Silicon Valley stereotypes we've grown accustomed to. It does not feature people who think that everything will be unrecognizably different tomorrow or the next day because of some restless tinkerer in his garage. It is neither a story about people who think technology will solve all our problems nor one about people who think technology is ineluctably bound to create apocalyptic new ones. It is not about disruption, at least not in the way that word tends to be used.

It is, in fact, three overlapping stories that converge in Google Translate's successful metamorphosis to A.I. — a technical story, an institutional story and a story about the evolution of ideas. The technical story is about one team on one product at one company, and the process by which they refined, tested and introduced a brand-new version of an old product in only about a quarter of the time

(the full article at https://nyti.ms/2hMtKOn is 39 pages long - highly recommended)

The New York Times | https://nyti.ms/2mF3loY

BOOK REVIEW | NONFICTION

Ray Kurzweil on How We'll End Up Merging With Our Technology

By RAY KURZWEIL MARCH 14, 2017

THINKING MACHINES
The Quest for Artificial Intelligence — and Where It's Taking Us Next
By Luke Dormehl
275 pp. TarcherPerigee. Paper, $16.

HEART OF THE MACHINE
Our Future in a World of Artificial Emotional Intelligence
By Richard Yonck
312 pp. Arcade Publishing. $25.99.

Books about science and especially computer science often suffer from one of two failure modes. Treatises by scientists sometimes fail to clearly communicate insights. Conversely, the work of journalists and other professional writers may exhibit a weak understanding of the science in the first place.

Luke Dormehl is the rare lay person — a journalist and filmmaker — who actually understands the science (and even the math) and is able to parse it in an edifying and exciting way. He is also a gifted storyteller who interweaves the personal stories with the broad history of artificial intelligence. I found myself turning the pages of "Thinking Machines" to find out what happens, even though I was there for much of it, and often in the very room.

Dormehl starts with the 1964 World's Fair — held only miles from where I lived as a high school student in Queens — evoking the anticipation of a nation working on sending a man to the moon. He identifies the early examples of artificial intelligence that captured my own excitement at the time, like IBM's demonstrations of automated handwriting recognition and language translation. He writes as if he had been there.

Dormehl describes the early bifurcation of the field into the Symbolic and Connectionist schools, and he captures key points that many historians miss, such as the uncanny confidence of Frank Rosenblatt, the Cornell professor who pioneered the first popular neural network (he called them "perceptrons"). I visited Rosenblatt in 1962 when I was 14, and he was indeed making fantastic claims for this technology, saying it would eventually perform a very wide range of tasks at human levels, including speech recognition, translation and even language comprehension. As Dormehl recounts, these claims were ridiculed at the time, and indeed the machine Rosenblatt showed me in 1962 couldn't perform any of these things. In 1969, funding for the neural net field was obliterated for about two decades when Marvin Minsky and his M.I.T. colleague Seymour Papert published the book "Perceptrons," which proved a theorem that perceptrons could not distinguish a connected figure (in which all parts are connected to each other) from a disconnected figure, something a human can do easily.

What Rosenblatt told me in 1962 was that the key to the perceptron achieving human levels of intelligence in many areas of learning was to stack the perceptrons in layers, with the output of one layer forming the input to the next. As it turns out, the Minsky-Papert perceptron theorem applies only to single-layer perceptrons. As Dormehl recounts, Rosenblatt died in 1971 without having had the chance to respond to Minsky and Papert's book. It would be decades before multi-layer neural nets proved Rosenblatt's prescience. Minsky was my mentor for 54 years until his death a year ago, and in recent years he lamented the "success" of his book and had become respectful of the recent gains in neural net technology. As Rosenblatt had predicted, neural nets were indeed providing near human-level (and in some cases superhuman levels) of performance on a wide range of intelligent tasks, from translating languages to driving cars to playing Go.

Dormehl examines the pending social and economic impact of artificial intelligence, for example on employment. He recounts the positive history of automation. In 1900, about 40 percent of American workers were employed on farms and over 20 percent in factories. By 2015, these figures had fallen to 2 percent on farms and 8.7 percent in factories. Yet for every job that was eliminated, we invented several new ones, with the work force growing from 24 million people (31 percent of the population in 1900) to 142 million (44 percent of the population in 2015). The average job today pays 11 times as much per hour in constant dollars as it did a century ago. Many economists are saying that while this may all be true, the future will be different because of the unprecedented acceleration of progress. Although expressing some cautions, Dormehl shares my optimism that we will be able to deploy artificial intelligence in the role of brain extenders to keep ahead of this economic curve. As he writes, "Barring some catastrophic risk, A.I. will represent an overall net positive for humanity when it comes to employment."

Many observers of A.I. and the other 21st-century exponential technologies like biotechnology and nanotechnology attempt to peer into the continuing accelerating gains and fall off the horse. Dormehl ends his book still in the saddle, discussing the prospect of conscious A.I.s that will demand and/or deserve rights, and the possibility of "uploading" our brains to the cloud. I recommend this book to anyone with a lay scientific background who wants to understand what I would argue is today's most important revolution, where it came from, how it works and what is on the horizon.

"Heart of the Machine," the futurist Richard Yonck's new book, contains its important insight in the title. People often think of feelings as secondary or as a sideshow to intellect, as if the essence of human intelligence is the ability to think logically. If that were true, then machines are already ahead of us. The superiority of human thinking lies in our ability to express a loving sentiment, to create and appreciate music, to get a joke. These are all examples of emotional intelligence, and emotion is at both the bottom and top of our thinking. We still have that old reptilian brain that provides our basic motivations for meeting our physical needs and to which we can trace feelings like anger and jealousy. The neocortex, a layer covering the brain, emerged in mammals two hundred million years ago and is organized as a

hierarchy of modules. Two million years ago, we got these big foreheads that house the frontal cortex and enabled us to process language and music.

Yonck provides a compelling and thorough history of the interaction between our emotional lives and our technology. He starts with the ability of the early hominids to fashion stone tools, perhaps the earliest example of technology. Remarkably the complex skills required were passed down from one generation to the next for over three million years, despite the fact that for most of this period, language had not yet been invented. Yonck makes a strong case that it was our early ability to communicate through pre-language emotional expressions that enabled the remarkable survival of this skill, and enabled technology to take root.

Yonck describes today's emerging technologies for understanding our emotions using images of facial expressions, intonation patterns, respiration, galvanic skin response and other signals — and how these instruments might be adopted by the military and interactive augmented reality experiences. And he recounts how all communication technologies from the first books to today's virtual reality have had significant sexual applications and will enhance sensual experiences in the future.

Yonck is a sure-footed guide and is not without a sense of humor. He imagines, for example, a scenario a few decades from now with a spirited exchange at the dinner table. "No daughter of mine is marrying a robot and that's final!" a father exclaims.

His daughter angrily replies: "Michael is a cybernetic person with the same rights you and I have! We're getting married and there's nothing you can do to change that!" She storms out of the room.

Yonck concludes that we will merge with our technology — a position I agree with — and that we have been doing so for a long time. He argues, as have I, that merging with future superintelligent A.I.s is our best strategy for ensuring a beneficial outcome. Achieving this requires creating technology that can understand and master human emotion. To those who would argue that such a quest is arrogantly playing God, he says simply: "This is what we do."

Ray Kurzweil, an inventor and futurist, is the author of “The Singularity Is Near” and “How to Create a Mind.”

A version of this review appears in print on March 19, 2017, on Page BR13 of the Sunday Book Review with the headline: Integrated Circuits.
