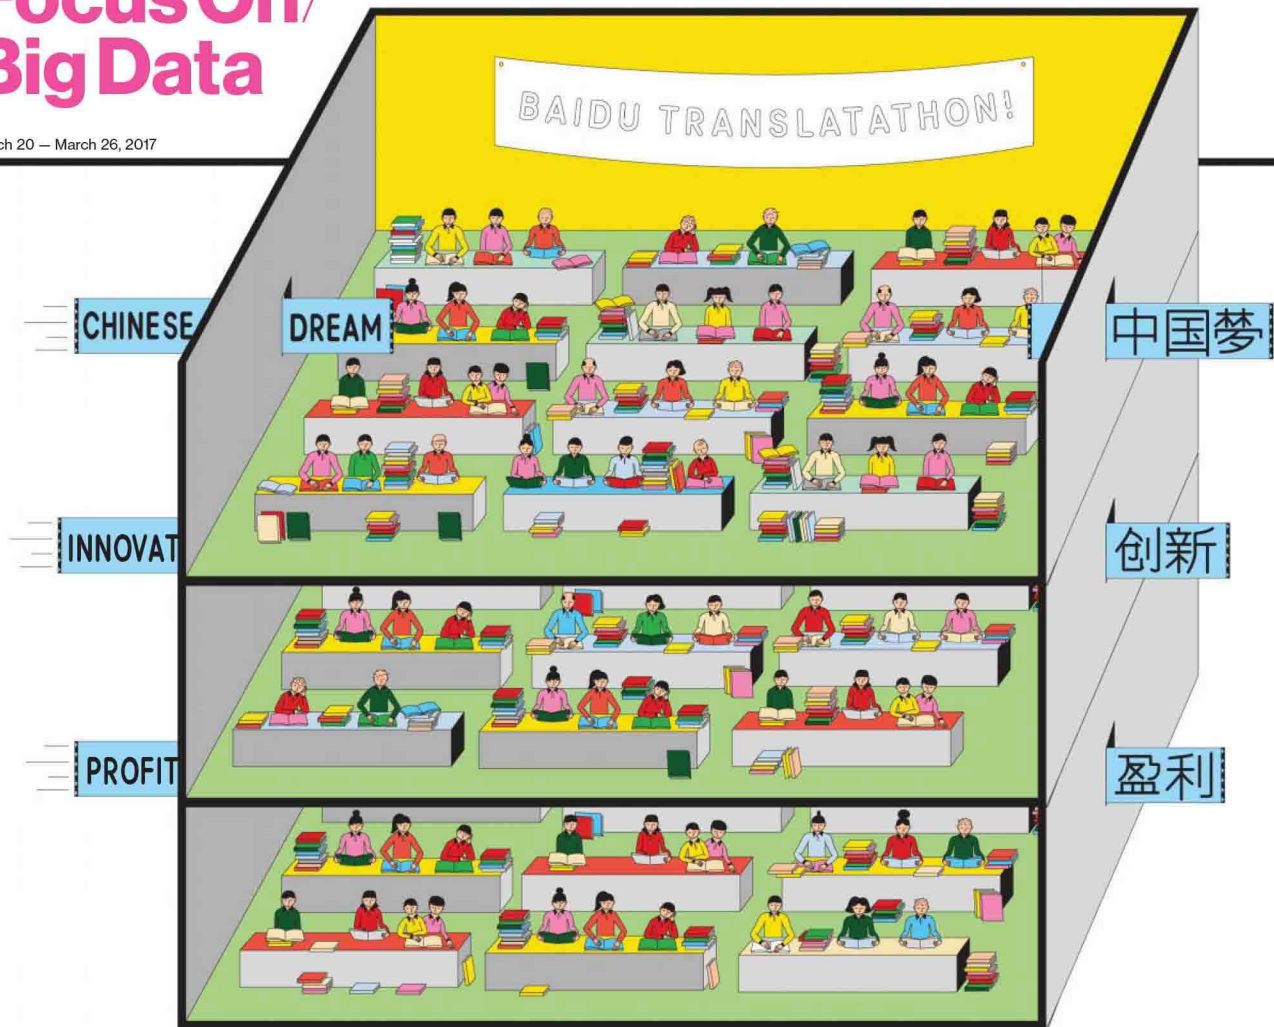


Focus On/ Big Data

March 20 — March 26, 2017



AI With Chinese Characteristics

► Baidu has a labor-intensive approach to deep learning

► “We literally dubbed it the Tchaikovsky Problem”

On Dec. 6, 2016, thousands of translators filed into office buildings across mainland China to pore over brochures, letters, and technical manuals, all in foreign languages, painstakingly rendering their texts in Chinese characters. This marathon carried on for 15 hours a day for an entire month. Clients that supplied the material received professional-grade Chinese versions of the originals at a bargain price. But **Baidu Inc.**, the Beijing-based company that organized the mass

translation, got something potentially more valuable: millions of English-Mandarin word pairs with which to train its online translation engine.

China is infamous for its knockoffs, whether luxury handbags or web startups. But the country's leadership seems to understand that when it comes to artificial intelligence, cheap imitations just won't do—not when its rivals include **Alphabet, Facebook, IBM,** and **Microsoft.** In February the National Development and Reform Commission appointed Baidu—often

described as the Google of China—to lead a new AI lab, signaling that Beijing believes the company has the makings of a national champion in this sphere.

Of the more than 20 billion yuan (\$2.9 billion) Baidu has spent on research and development over the past two and a half years, most has been on AI, according to comments co-founder and Chief Executive Officer Robin Li made at the lab's launch last month. But China's national interest isn't his main motivation: Baidu's revenue growth fell to about 6 percent

last year, from an average of more than 30 percent over the prior three years. The search ad business, which contributed the lion's share of its 70.5 billion yuan in sales in the fiscal year ended on Dec. 31, is under siege from local rivals. A September report from EMarketer Inc. noted that **Alibaba Group Holding Ltd.** had overtaken Baidu to become the leader in China's digital ad market. Baidu hopes AI can help it reclaim share in search, as well as ensure success in newer ventures.

That's key as the 17-year-old company's attempts at diversification have produced mixed results. The number of daily visitors to its group buying site, Nuomi, dropped 59 percent in the 12 months through February 2017; its Waimai food delivery service lags in third place, according to Natalie Wu, an analyst with China International Capital Corp. The Netflix-like streaming video service iQiyi.com is hugely popular, but it will take 12 billion yuan to keep it stocked with content this year, estimates analyst Ella Ji with China Renaissance Securities (Hong Kong) Ltd.

Those faltering efforts mean Baidu's push into AI is taking on greater importance. "The era of mobile internet has ended," said Li in a March 10 interview. "We're going to aggressively invest in AI, and I think it's going to benefit a lot of people and transform industry after industry."

In January the company named former Microsoft Corp. executive Qi Lu as its chief operating officer, with a mandate to reshape the company around such technologies as deep learning, augmented reality, and image recognition. He joins Chief Scientist Andrew Ng, a Stanford academic who worked on Alphabet Inc.'s deep learning group before decamping to Baidu in 2014. Under Ng's watch, the company's AI team, which is scattered across research labs in Beijing, Shenzhen, Shanghai, and Sunnyvale, Calif., has grown to 1,300 and is expected to increase by several hundred more hires this year. "A ton of stuff is invented in China, and a ton of stuff is invented in the U.S.," says Ng, who's based in Silicon Valley. "By having

people in both countries, we see the latest trends."

On the day in May 2014 that the Sunnyvale research center opened, Ng and his top lieutenant, Adam Coates, sat down in front of a blank whiteboard to identify their first project. After drawing up a list of possibilities (and challenges), they settled on speech recognition as a foundation on which they could build a series of other offerings.

By mid-2015, the 50-person team had a product called Deep Speech that could decipher much of what was said in English. Rather than picking apart phrases word by word, the software parsed through vast reams of language data and then extrapolated patterns, a process known as deep learning. The system could transcribe speech more accurately than traditional engines that rely on vocabulary lists and phonetic dictionaries, Ng says, because it took into account a word's context to determine its meaning.

One thing that consistently tripped it up, though, were words and names that had over time crept into the English lexicon from other languages. "If you want to say 'Play music by Tchaikovsky,' the software would return answers like 'Play music and try cough ski,'" says Coates, whom Ng recruited from Stanford. "We literally dubbed it the Tchaikovsky Problem."

Instead of simply adding "Tchaikovsky" to the system's vocabulary list, Baidu's programmers had to help Deep Speech teach itself to understand the word. That involved pumping in even more data to help the system put things in context.

Shiqi Zhao, the Beijing-based associate director of Baidu's natural language processing department, recalls that as a computer science student at China's Harbin Institute of Technology he had only 2 million word pairs of English-to-Chinese terms to play with while working on computer-based translation; Baidu has about 100 million. However, that's still far fewer than Alphabet's

500 million, according to a 2016 article in *Science* magazine that featured one of the U.S. company's research scientists, Quoc V. Le.

To help close the gap, Baidu has resorted to an age-old tactic: throw lots of people at the problem. The company now facilitates manual translations year-round and stages marathon events such as the one in December at regular intervals, in which it offers clients prizes such as smartphones and water purifiers. The data collected help enhance the performance of its Baidu Translate engine as well as further the development of Deep Speech.

The software created by the Sunnyvale team had its commercial debut in July 2016 with the release of TalkType, a keyboard app with a talk-to-text feature. The technology has since been incorporated into other products, including a Siri-like personal assistant named DuMi in China and DuEr everywhere else. (DuMi is a fusion of "du" from Baidu and "mi," which means "secretary" in Mandarin; DuEr sounds like "doer.") The machine learning Baidu has inculcated into Deep Speech is helping it animate other products with intelligence. For instance, it's the secret sauce in Xiaoyu Zaijia ("Little Fish"), a voice-controlled robot, à la Amazon Echo, that Baidu showed off at the CES show in Las Vegas in January.

Baidu's portfolio of web properties gives it access to one of the largest and most detailed sets of consumer data ever produced in China, which—in theory at least—should give it an edge in building AI-infused products and services for the mainland. Thanks to Nuomi and Waimai, the company knows what Chinese households

buy and eat, while Ctrip.com, the world's second-largest online travel agent, reveals where they want to holiday. Every month 665 million smartphone users surf its mobile portal and apps, while 341 million use Baidu Maps to reach their destination.

"It's a mistake to think of AI as a product—it underpins and enables product," says HSBC ▶

◀ Holdings Plc analyst Chi Tsang. "Think of all the use cases."

The new AI products aren't contributing much to Baidu's bottom line yet. But the company's nascent expertise in this area could help it achieve dominance in segments where it's already present and propel it into new ones, such as cloud computing and self-driving cars. "In the next three to five years all those areas have the potential to become another Baidu," company President Zhang Ya-Qin says, referring to Baidu's \$60.2 billion market capitalization. "Right now it's time to make some bets."

—David Ramli and Alex Webb

The bottom line Baidu's 1,300-person AI team is writing software to improve everything from translation to food ordering.



Baidu showed off Little Fish, a voice-activated robot, at CES