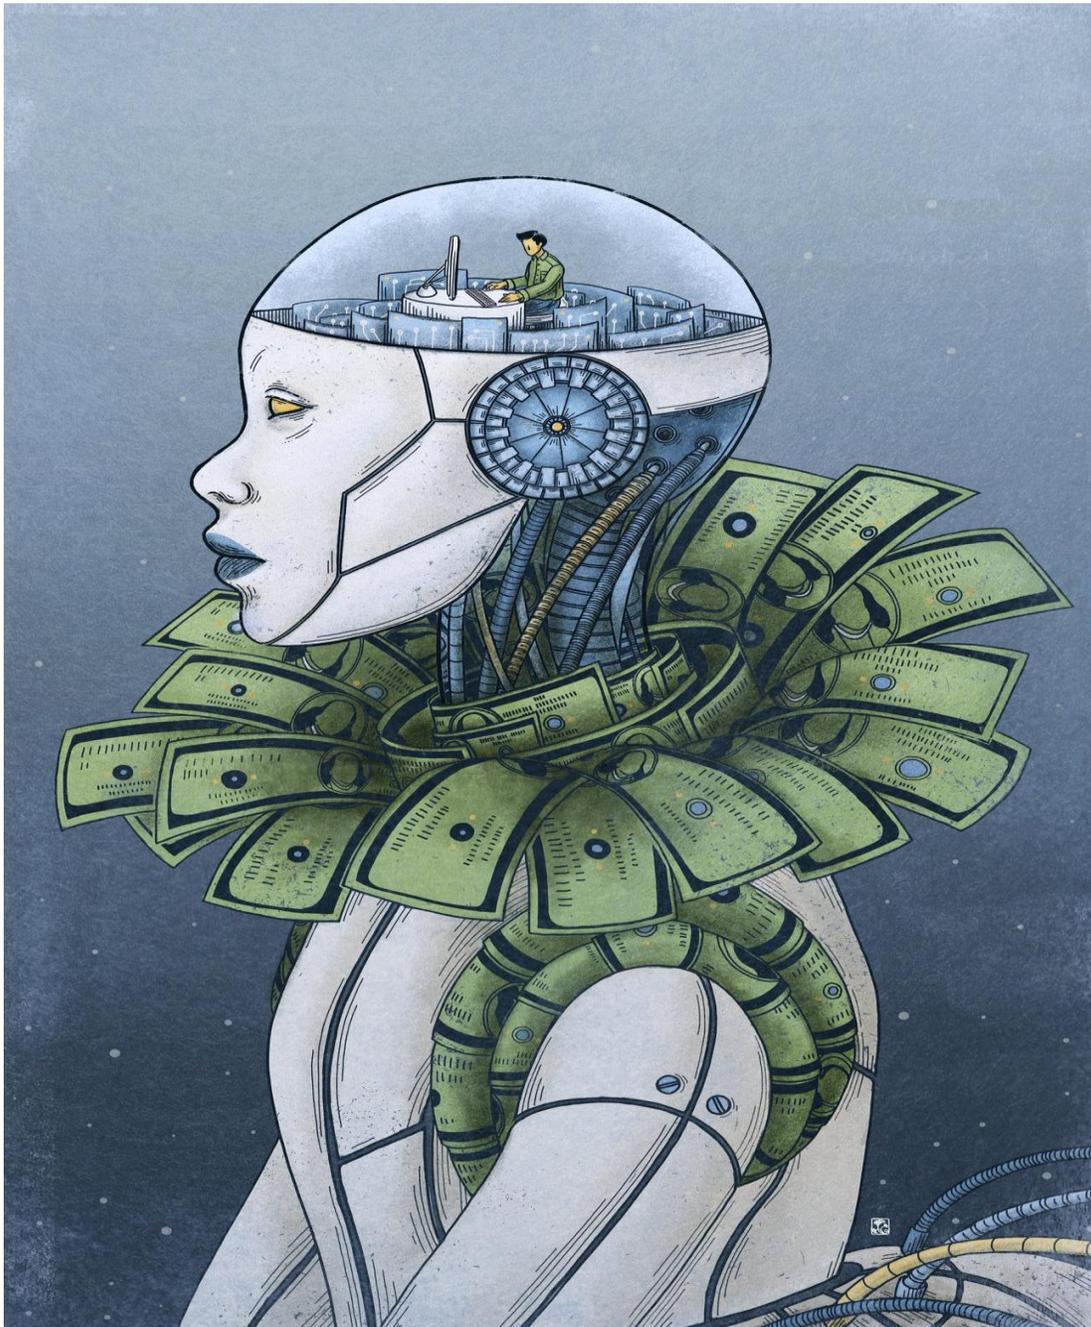


Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent

Nearly all big tech companies have an artificial intelligence project, and they are willing to pay experts millions of dollars to help get it done.

By CADE METZ OCT. 22, 2017



SAN FRANCISCO — Silicon Valley's start-ups have always had a recruiting advantage over the industry's giants: Take a chance on us and we'll give you an ownership stake that could make you rich if the company is successful.

Now the tech industry's race to embrace artificial intelligence may render that advantage moot — at least for the few prospective employees who know a lot about A.I.

Tech's biggest companies are placing huge bets on artificial intelligence, banking on things ranging from face-scanning smartphones and conversational coffee-table gadgets to computerized health care and autonomous vehicles. As they chase this future, they are doling out salaries that are startling even in an industry that has never been shy about lavishing a fortune on its top talent.

Typical A.I. specialists, including both Ph.D.s fresh out of school and people with less education and just a few years of experience, can be paid from \$300,000 to \$500,000 a year or more in salary and company stock, according to nine people who work for major tech companies or have entertained job offers from them. All of them requested anonymity because they did not want to damage their professional prospects.

Well-known names in the A.I. field have received compensation in salary and shares in a company's stock that total single- or double-digit millions over a four- or five- year period. And at some point they renew or negotiate a new contract, much like a professional athlete.

At the top end are executives with experience managing A.I. projects. In a court filing this year, Google revealed that one of the leaders of its self-driving-car division, Anthony Levandowski, a longtime employee who started with Google in 2007, took home over \$120 million in incentives before joining Uber last year through the acquisition of a start-up he had co-founded that drew the two companies into a court fight over intellectual property.

Salaries are spiraling so fast that some joke the tech industry needs a National Football League-style salary cap on A.I. specialists. "That would make things easier," said Christopher Fernandez, one of Microsoft's hiring managers. "A lot easier."

There are a few catalysts for the huge salaries. The auto industry is competing with Silicon Valley for the same experts who can help build self-driving cars. Giant tech companies like Facebook and Google also have plenty of money to throw around and problems that they think A.I. can help solve, like building digital assistants for smartphones and home gadgets and spotting offensive content.

Most of all, there is a shortage of talent, and the big companies are trying to land as much of it as they can. Solving tough A.I. problems is not like building the flavor- of-the-month smartphone app. In the entire world, fewer than 10,000 people have the skills necessary to tackle serious artificial intelligence research, according to Element AI, an independent lab in Montreal.

"What we're seeing is not necessarily good for society, but it is rational behavior by these companies," said Andrew Moore, the dean of computer science at Carnegie Mellon University, who previously worked at Google. "They are anxious to ensure that they've got this small cohort of people" who can work on this technology.

Costs at an A.I. lab called DeepMind, acquired by Google for a reported \$650 million in 2014, when it employed about 50 people, illustrate the issue. Last year, according to the company's recently released annual financial accounts in Britain, the lab's "staff costs" as it expanded to 400 employees totaled \$138 million. That comes out to \$345,000 an employee.

"It is hard to compete with that, especially if you are one of the smaller companies," said Jessica Cataneo, an executive recruiter at the tech recruiting firm CyberCoders.

The cutting edge of artificial intelligence research is based on a set of mathematical techniques called deep neural networks. These networks are mathematical algorithms that can learn tasks on their own by analyzing data. By looking for patterns in millions of dog photos, for example, a neural network can learn to recognize a dog. This mathematical idea dates back to the 1950s, but it remained on the fringes of academia and industry until about five years ago.

By 2013, Google, Facebook and a few other companies started to recruit the relatively few researchers who specialized in these techniques. Neural networks now help recognize faces in photos posted to Facebook, identify commands spoken into living-room digital assistants like the Amazon Echo and instantly translate foreign languages on Microsoft's Skype phone service.

Using the same mathematical techniques, researchers are improving self-driving cars and developing hospital services that can identify illness and disease in medical scans, digital assistants that can not only recognize spoken words but understand them, automated stock-trading systems and robots that pick up objects they've never seen before.

With so few A.I. specialists available, big tech companies are also hiring the best and brightest of academia. In the process, they are limiting the number of professors who can teach the technology.

Uber hired 40 people from Carnegie Mellon's groundbreaking A.I. program in 2015 to work on its self-driving-car project. Over the last several years, four of the best-known A.I. researchers in academia have left or taken leave from their

professorships at Stanford University. At the University of Washington, six of 20 artificial intelligence professors are now on leave or partial leave and working for outside companies.

"There is a giant sucking sound of academics going into industry," said Oren Etzioni, who is on leave from his position as a professor at the University of Washington to oversee the nonprofit Allen Institute for Artificial Intelligence.

Some professors are finding a way to compromise. Luke Zettlemoyer of the University of Washington turned down a position at a Google-run Seattle laboratory that he said would have paid him more than three times his current salary (about \$180,000, according to public records). Instead, he chose a post at the Allen Institute that allowed him to continue teaching.

"There are plenty of faculty that do this, splitting their time in various percentages between industry and academia," Mr. Zettlemoyer said. "The salaries are so much higher in industry, people only do this because they really care about being an academian."

To bring in new A.I. engineers, companies like Google and Facebook are running classes that aim to teach “deep learning” and related techniques to existing employees. And nonprofits like Fast.ai and companies like Deeplearning.ai, founded by a former Stanford professor who helped create the Google Brain lab, offer online courses.

The basic concepts of deep learning are not hard to grasp, requiring little more than high-school-level math. But real expertise requires more significant math and an intuitive talent that some call “a dark art.” Specific knowledge is needed for fields like self-driving cars, robotics and health care.

In order to keep pace, smaller companies are looking for talent in unusual places. Some are hiring physicists and astronomers who have the necessary math skills. Other start-ups from the United States are looking for workers in Asia, Eastern Europe and other locations where wages are lower.