AI and Economic Productivity: Expect Evolution, Not Revolution

spectrum.ieee.org/computing/software/ai-and-economic-productivity-expect-evolution-not-revolution

Jeffrey Funk, *spectrum.ieee*, December 5, 2019

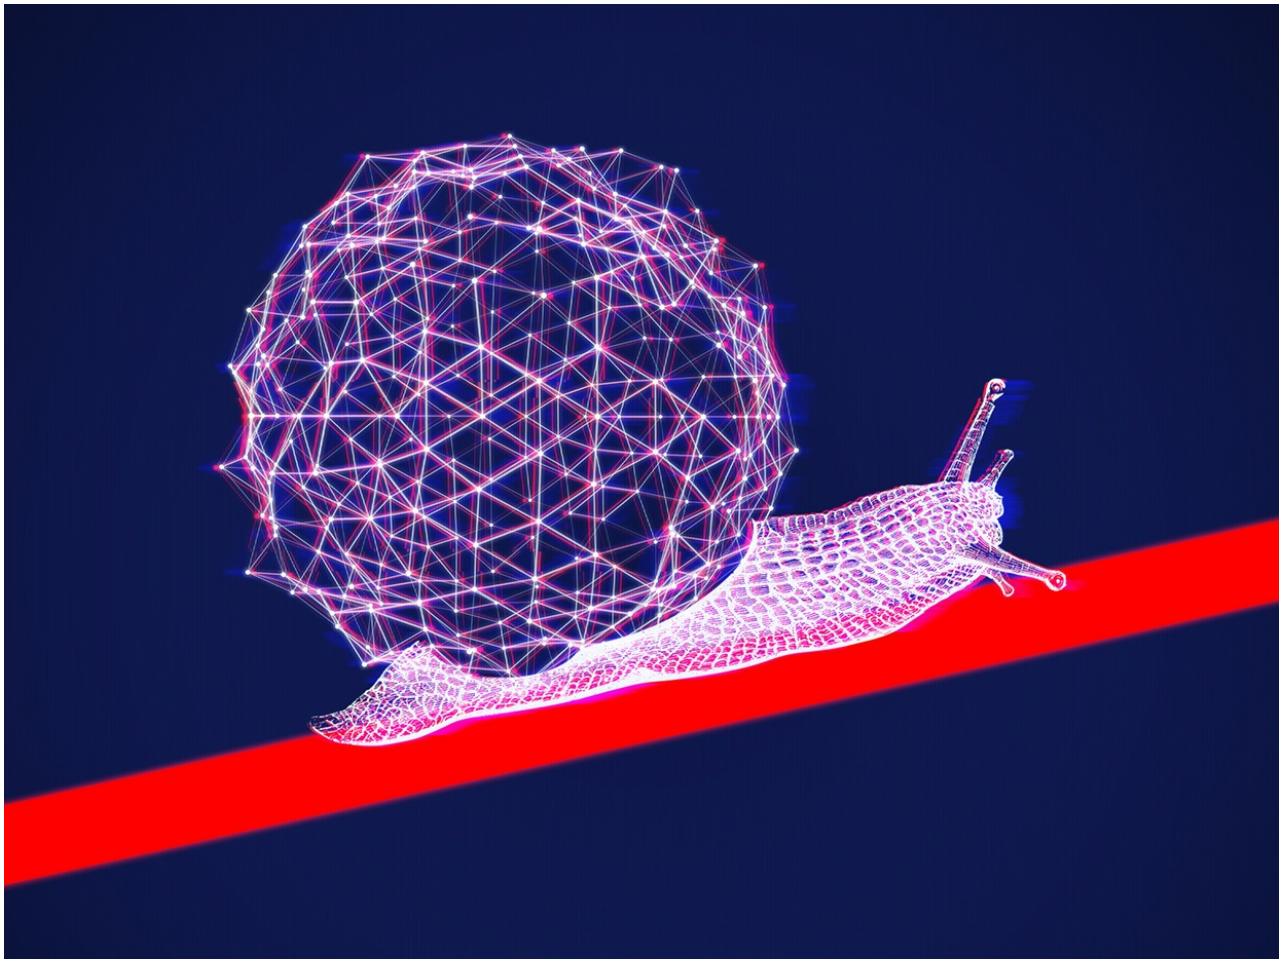

Illustration: Edmon de Haro

In 2016, London-based DeepMind Technologies, a subsidiary of Alphabet (which is also the parent company of Google), startled industry watchers when it reported that the application of artificial intelligence had reduced the cooling bill at a Google data center by a whopping 40 percent. What's more, we learned that year, DeepMind was starting to work with National Grid in the United Kingdom to save energy throughout the country using deep learning to optimize the flow of electricity.

Could AI really slash energy usage so profoundly? In the three years that have passed, I've searched for articles on the application of AI to other data centers but find no evidence of important gains. What's more, DeepMind's talks with National Grid about energy have broken down. And the financial results for DeepMind certainly don't suggest that customers are lining up for its services: For 2018, the company reported losses of US $571 million on revenues of $125 million, up from losses of $366 million in 2017. Last April, *The Economist* characterized DeepMind's 2016 announcement as a publicity stunt, quoting one inside source as saying, "[DeepMind just wants] to have some PR so they can claim some value added within Alphabet."

This episode encouraged me to look more deeply into the economic promise of AI and the rosy projections made by champions of this technology within the financial sector. This investigation was just the latest twist on a long- standing interest of mine. In the early 1980s, I wrote a doctoral dissertation on the economics of robotics and AI, and throughout my career as a professor and technology consultant I have followed the economic projections for AI, including detailed assessments by consulting organizations such as Accenture, PricewaterhouseCoopers International (PwC), and McKinsey.

These analysts have lately been asserting that AI-enabled technologies will dramatically increase economic output. Accenture claims that by 2035 AI will double growth rates for 12 developed countries and increase labor productivity by as much as a third. PwC claims that AI will add $15.7 trillion to the global economy by 2030, while McKinsey projects a $13 trillion boost by that time.

Other forecasts have focused on specific sectors such as retail, energy, education, and manufacturing. In particular, the McKinsey Global Institute assessed the impact of AI on these four sectors in a 2017 report titled *Artificial Intelligence: The New Digital Frontier?* and did so for a much longer list of sectors in a 2018 report. In the latter, the institute concluded that AI techniques "have the potential to create between $3.5 trillion and $5.8 trillion in value annually across nine business functions in 19 industries. This constitutes about 40 percent of the overall $9.5 trillion to $15.4 trillion annual impact that could potentially be enabled by all analytical techniques."

Wow. These are big numbers. If true, they create a powerful incentive for companies to pursue AI—with or without help from McKinsey consultants. But are these predictions really valid?

Many of McKinsey's estimates were made by extrapolating from claims made by various startups. For instance, its prediction of a 10 percent improvement in energy efficiency in the U.K. and elsewhere was based on the purported success of DeepMind and also of Nest Labs, which became part of Google's hardware division in 2018. In 2017, Nest, which makes a smart thermostat and other intelligent products for the home, lost $621 million on revenues of $726 million. That fact doesn't mesh with the notion that Nest and similar companies are contributing, or are poised to contribute, hugely to the world economy.

So I decided to investigate more systematically how well such AI startups were doing. I found that many were proving not nearly as valuable to society as all the hype would suggest. This assertion will certainly rub a lot of people the wrong way, the analysts at McKinsey among them. So I'd like to describe here how I reached my much more pessimistic conclusions.

My investigation of Nest Labs expanded into a search for evidence that smart meters in general are leading to large gains in energy efficiency. In 2016, the British government began a coordinated campaign to install smart meters throughout the country by 2020. And since 2010, the U.S. Department of Energy has invested some $4.5 billion installing more than 15 million smart meters throughout the United States. Curiously enough, all

that effort has had little observed impact on energy usage. The U.K. government recently revised downward the amount it figures a smart meter will save each household annually, from £26 to just £11. And the cost of smart meters and their installation has risen, warns the U.K.'s National Audit Office. All of this is not good news for startups banking on the notion that smart thermostats, smart home appliances, and smart meters will lead to great energy savings.

Are other kinds of AI startups having a greater positive effect on the economy? Tech sector analyst CB Insights reports that overall venture capital funding in the United States was $115 billion in 2018 [PDF], of which $9.3 billion went to AI startups. While that's just 8 percent of the total, it's still a lot of money, indicating that there are many U.S. startups working on AI (although some overstate the role of AI in their business plans to acquire funding).

To probe further, I gathered data on the U.S. AI startups that have received the most funding and looked at which industries they were hoping to disrupt. The reason for focusing on the United States is that it has the longest history of startup success, so it seems likely that its AI startups are more apt to flourish than those in other countries. My intention was to evaluate whether these U.S. startups had succeeded in shaking up various industries and boosting productivity or whether they promise to do so shortly.

U.S. AI Startups With More Than $100 Million in Equity Funding

Startup name	Equity funding unless otherwise stated	Year founded	Type of product or service
Tanium	$6.5 billion valuation	2007	Cybersecurity
Indigo Agriculture	$3.5 billion valuation	2016	Microbial seed treatments
CrowdStrike	$3.4 billion valuation	2011	Security
Automation Anywhere	$2.6 billion valuation	2003	Robotic process automation
Avant	$1.9 billion valuation	2002	Credit scores for personal lending
Uptake Technologies	$2.3 billion valuation	2014	Internet of Things platform
Flatiron Health	$1.9 billion acquisition (by Roche in 2018)	2012	Health care

Xant	$1.7 billion valuation	2004	Platform for sales team
Cylance	$1.4 billion acquisition (by BlackBerry in 2018)	2012	Security
Tempus Labs	$1.1 billion valuation	2015	Health care
ZipRecruiter	$1.0 billion valuation	2010	Human resources
OpenAI	$1.0 billion valuation	2015	Creates software that benefits “all of humanity”
Zoox	$800 million	2014	Driverless vehicles
Dataminr	$577 million	2009	Cybersecurity
Zymergen	$574 million	2013	Chemical fermentation
UiPath	$409 million	2005	Robotic process automation
ZestFinance	$268 million	2009	Financial technology and insurance
Cybereason	$188 million	2012	Cybersecurity
Nauto	$183 million	2015	In-vehicle alerts to reduce distracted driving
Sentient Technologies	$174 million	2007	Parallel AI processing for large, complex data sets
CloudMinds	$130 million	2015	Cloud robotics
Brain Corp.	$125 million	2009	Converts manual equipment into intelligent robots
Afiniti	$122 million	2005	Analyzes patterns of human interaction
Vicarious	$122 million	2010	Interprets photos and videos; solves CAPTCHA
WorkFusion	$121 million	2010	Robotic process automation

Wave Computing	$117 million	2008	AI chips
SoundHound	$115 million	2005	Music-recognition software
DataRobot	$111 million	2012	Automates data-science work
Petuum	$108 million	2016	Parallel machine learning for complex data sets
Shape Security	$106 million	2011	Cybersecurity
Ayasdi	$106 million	2008	Software platform
Endgame	$96 million	2008	Cybersecurity
Nuro	$92 million	2016	Small driverless delivery vehicles
Conversica	$87 million	2007	Sales and marketing
Upstart	$86 million	2012	Financial technology and insurance
Freenome	$79 million	2014	Cancer screening
Orbital Insight	$79 million	2013	Geospatial analytics
Trifacta	$76 million	2012	Data preparation before processing
Algolia	$74 million	2012	Custom Web search tools for businesses
H2O	$74 million	2011	Open-source platform

In all, I examined 40 U.S. startups working on AI. These either had valuations greater than $1 billion or had more than $70 million in equity funding. Other than two that had been acquired by public companies, the startups I looked at are all private firms. I found their names and product offerings in lists of leading startups that Crunchbase, Fortune, and Datamation had compiled and published. I then updated my data set with more recent news about these companies (including reports of some shutdowns).

I categorized these 40 startups by the type of product or service they offered. Seventeen are working on what I would call basic computer hardware and software (Wave Computing and OpenAI, respectively, are examples), including cybersecurity (CrowdStrike, for instance). That is, I included in this category companies building tools that are intended to support the computing environment itself.

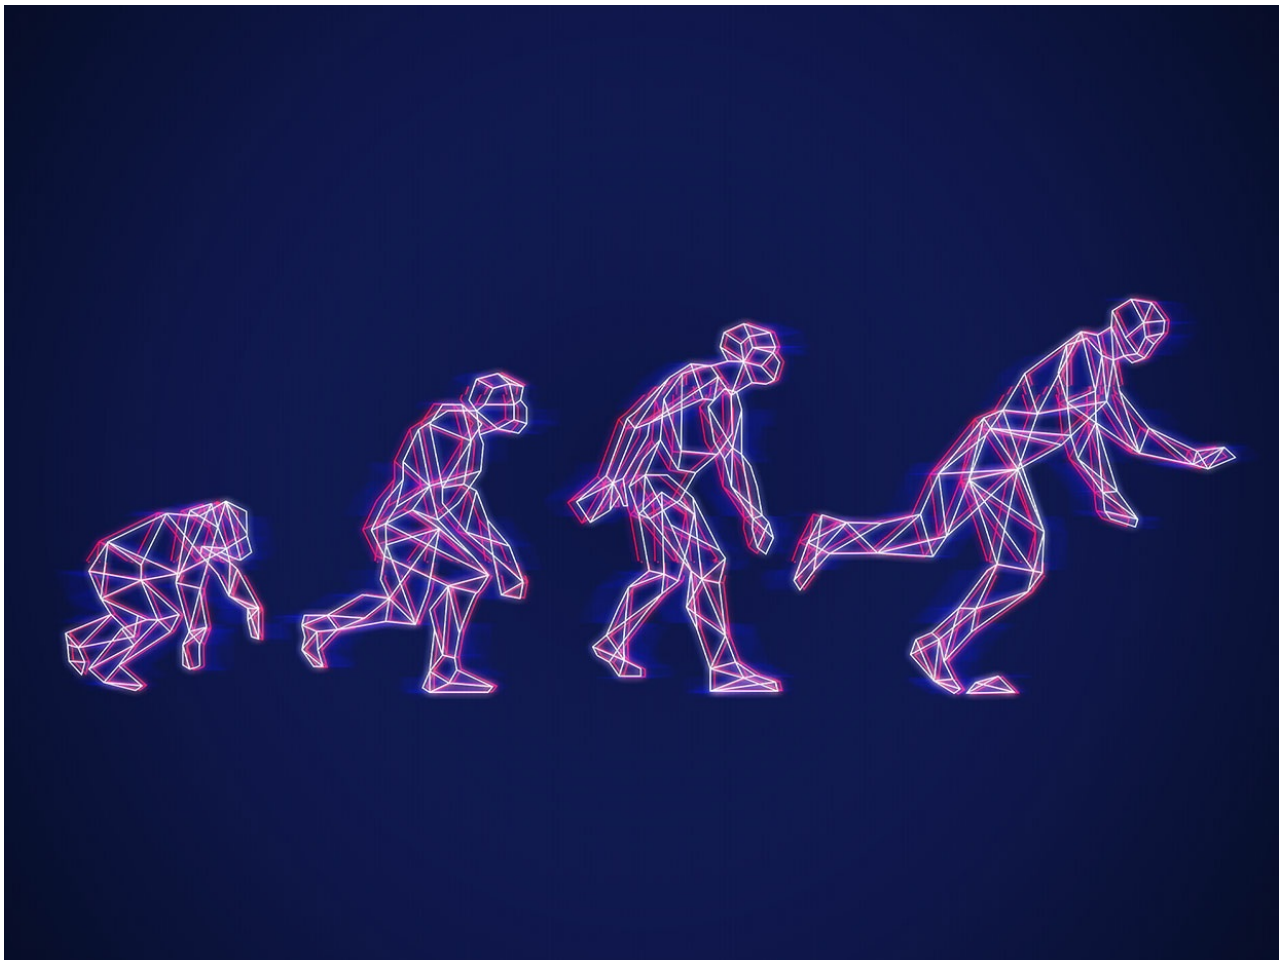

Illustration: Edmon de Haro

Making up another large fraction—8 of the 40—are companies that develop software that automates various tasks. The robotic process-automation software being developed by Automation Anywhere, UiPath, and WorkFusion, for example, enables higher productivity among professionals and other white-collar workers. Software from Brain Corp. converts manual equipment into intelligent robots. Algolia, Conversica, and Xant offer software to improve sales and marketing. ZipRecruiter targets human resources.

The remaining startups on my list are spread among various industries. Three (Flatiron Health, Freenome, Tempus Labs) work in health care; three more (Avant, Upstart, ZestFinance) are focused on financial technology; two (Indigo Agriculture, Zymergen) target agriculture or synthetic biology; and three others (Nauto, Nuro, Zoox) involve transportation. There is just one startup each for geospatial analytics (Orbital Insight), patterns of human interaction (Afiniti), photo/video recognition (Vicarious), and music recognition (SoundHound).

Are there indications that these startups will bring large productivity improvements in the near future? In my view, software that automates tasks normally carried out by white-collar workers is probably the most promising of the products and services that AI is being applied to. Similar to past improvements in tools for white-collar professionals, including Excel for accountants and computer-aided design for engineers and architects, these types of AI-based tools have the greatest potential impact on productivity. For instance, there are high hopes for generative design, in which teams of people input constraints and the system proposes specific designs.

But looking at the eight startups on my list that are working on automation tools for white-collar workers, I realized that they are not targeting things that would lead to much higher productivity. Three of them are focused on sales and marketing, which is often a zero-sum game: The company with the best software takes customers from competitors, with only small increases in productivity under certain conditions. Another one of these eight companies is working on human-resource software, whose productivity benefits may be larger than those for sales and marketing but probably not as large as you'd get from improved robotic process automation.

This leaves four startups that do offer such software, which may lead to higher productivity and lower costs. But even among these startups, none currently offers software that helps engineers and architects become more productive through, for example, generative design. Software of this kind isn't coming from the largest startups, perhaps because there is a strong incumbent, Autodesk, or because the relevant AI is still not developed enough to provide truly useful tools in this area.

The relatively large number of startups I classified as working on basic hardware and software for computing (17) also suggests that productivity improvements are still many years away. Although basic hardware and software are a necessary part of developing higher-level AI-based tools, particularly ones utilizing machine learning, it will take time for the former to enable the latter. I suppose this situation simply reflects that AI is still in its infancy. You certainly get that impression from companies like OpenAI: Although it has received $1 billion in funding (and a great deal of attention), the vagueness of its mission—"Benefiting all of humanity"—suggests that it will take many years yet for specific useful products and services to evolve from this company's research.

The large number of these startups that are focused on cybersecurity (seven) highlights the increasing threat of security problems, which raise the cost of doing business over the Internet. AI's ability to address cybersecurity issues will likely make the Internet more safe, secure, and useful. But in the end, this thrust reflects yet higher costs in the future for Internet businesses and will not, to my mind, lead to large productivity improvements within the economy as a whole.

If not from the better software tools it brings, where will AI bring substantial economic gains? Health care, you would think, might benefit greatly from AI. Yet the number of startups on my list that are applying AI to health care (three) seems oddly small if that

were really the case. Perhaps this has something to do with IBM's experience with its Watson AI, which proved a disappointment when it was applied to medicine.

Still, many people remain hopeful that AI-fueled health care startups will fill the gap left by Watson's failures. Arguing against this is Robert Wachter, who points out that it's much more difficult to apply computers to health care than to other sectors. His 2015 book, *The Digital Doctor: Hope, Hype, and Harm at the Dawn of Medicine's Computer Age*, details the many reasons that health care lags other industries in the application of computers and software. It's not clear that adding AI to the mix of digital technologies available will do anything to change the situation.

There are also some big applications missing from the list of well-funded AI startups. Housing represents the largest category of consumer expenditures in the United States, but none of these startups are addressing this sector of the economy at all. Transportation is the second largest expenditure, and it is the focus of just three of these startups. One is working on a product that identifies distracted drivers. Another intends to provide automated local deliveries. Only one startup on the list is developing driverless passenger vehicles. That there is only one working on self-driving cars is consistent with the pessimism recently expressed by executives of Ford, General Motors, and Mercedes-Benz about the prospects for driverless vehicles taking to the streets in large numbers anytime soon, even though $35 billion has already been spent on R&D for them.

Admittedly, my assessment of what these 40 companies are doing and whether their offerings will shake up the world over the next decade is subjective. Perhaps it makes better sense to consider a more objective measure of whether these companies are providing value to the world economy: their profitability.

Alas, good financial data is not available on privately held startups, and only two of the companies on my list are now part of public companies. Also, startups often take years to turn a profit (Amazon took seven years). So there isn't a lot to go on here. Still, there are some broad trends in the tech sector that are quite telling.

The fraction of tech companies that are profitable by the time they go public dropped from 76 percent in 1980 to just 17 percent in 2018, even though the average time to IPO has been rising—it went from 2.8 years in 1998 to 7.7 years in 2016, for example. Also, the losses of some well-known startups that took a long time to go public are huge. For instance, none of the big ride-sharing companies are making a profit, including those in the United States (Uber and Lyft), China, India, and Singapore, with total losses of about $5 billion in 2018. Most bicycle and scooter sharing, office sharing, food delivery, P2P (peer-to peer) lending, health care insurance and analysis, and other consumer service startups are also losing vast amounts of money, not only in the United States but also in China and India.

Most of the 40 AI startups I examined will probably stay private, at least in the near term. But even if some do go public several years down the road, it's unlikely they'll be profitable at that point, if the experience of many other tech companies is any guide. It may take these companies years more to achieve the distinction of making more money than they are spending.

For the reasons I've given, it's very hard for me to feel confident that any of the AI startups I examined will provide the U.S. economy with a big boost over the next decade. Similar pessimism is also starting to emerge from such normally cheery publications as *Technology Review* and *Scientific American.* Even the AI community is beginning to express concerns in books such as *The AI Delusion* and *Rebooting AI: Building Artificial Intelligence We Can Trust*, concerns that are growing amid the rising hype about many new technologies.

The most promising areas for rapid gains in productivity are likely to be found in robotic process automation for white-collar workers, continuing a trend that has existed for decades. But these improvements will be gradual, just as those for computer-aided design and computer-aided engineering software, spreadsheets, and word processing have been.

Viewed over the span of decades, the value of such software is impressive, bringing huge gains in productivity for engineers, accountants, lawyers, architects, journalists, and others—gains that enabled some of these professionals (particularly engineers) to enrich the global economy in countless ways.

Such advances will no doubt continue with the aid of machine learning and other forms of AI. But they are unlikely to be nearly as disruptive—for companies, for workers, or for the economy as a whole—as many observers have been arguing.

Updated 7 February 2020.

About the Author

Jeffrey Funk retired from the National University of Singapore in 2017, where he taught (among other subjects) a course on the economics of new technology as a professor of technology management. He remains based in Singapore, where he consults in various areas of technology and business.