



[Artificial Intelligence](#)

AI is reinventing the way we invent

The biggest impact of artificial intelligence will be to help humans make discoveries we couldn't make on our own.

by **David Rotman**

Feb 15, 2019

Regina Barzilay's office at MIT affords a clear view of the Novartis Institutes for Biomedical Research. Amgen's drug discovery group is a few blocks beyond that. Until recently, Barzilay, one of the world's leading researchers in artificial intelligence, hadn't given much thought to these nearby buildings full of chemists and biologists. But as AI and machine learning began to perform ever more impressive

feats in image recognition and language comprehension, she began to wonder: could it also transform the task of finding new drugs?

The problem is that human researchers can explore only a tiny slice of what is possible. It's estimated that there are as many as 10^{60} potentially drug-like molecules—more than the number of atoms in the solar system. But traversing seemingly unlimited possibilities is what machine learning is good at. Trained on large databases of existing molecules and their properties, the programs can explore all possible related molecules.

Drug discovery is a hugely expensive and often frustrating process. Medicinal chemists must guess which compounds might make good medicines, using their knowledge of how a molecule's structure affects its properties. They synthesize and test countless variants, and most are failures. "Coming up with new molecules is still an art, because you have such a huge space of possibilities," says Barzilay. "It takes a long time to find good drug candidates."

By speeding up this critical step, deep learning could offer far more opportunities for chemists to pursue, making drug discovery much quicker. One advantage: machine learning's often quirky imagination. "Maybe it will go in a different direction that a human wouldn't go in," says Angel Guzman-Perez, a drug researcher at Amgen who is working with Barzilay. "It thinks differently."

Others are using machine learning to try to invent new materials for clean-tech applications. Among the items on the wish list are improved batteries for storing power on the electric grid and organic solar cells, which could be far cheaper to make than today's bulky silicon-based ones.

Such breakthroughs have become harder and more expensive to attain as chemistry, materials science, and drug discovery have grown mind-bogglingly complex and saturated with data. Even as the pharmaceutical and biotech industries pour money into research, the number of new drugs based on novel molecules has been flat over the last few decades. And we're still stuck with lithium-ion batteries that date to the early 1990s and designs for silicon solar cells that are also decades old.

The complexity that has slowed progress in these fields is where deep learning excels. Searching through multidimensional space to come up with valuable predictions is "AI's sweet spot," says Ajay Agrawal, an economist at the Rotman School of

Management in Toronto and author of the best-selling *Prediction Machines: The Simple Economics of Artificial Intelligence*.

In a recent paper, economists at MIT, Harvard, and Boston University argued that AI's greatest economic impact could come from its potential as a new "method of invention" that ultimately reshapes "the nature of the innovation process and the organization of R&D."

Iain Cockburn, a BU economist and coauthor of the paper, says: "New methods of invention with wide applications don't come by very often, and if our guess is right, AI could dramatically change the cost of doing R&D in many different fields." Much of innovation involves making predictions based on data. In such tasks, Cockburn adds, "machine learning could be much faster and cheaper by orders of magnitude."

In other words, AI's chief legacy might not be driverless cars or image search or even Alexa's ability to take orders, but its ability to come up with new ideas to fuel innovation itself.

Ideas are getting expensive

Late last year, Paul Romer won the economics Nobel Prize for work done during the late 1980s and early 1990s that showed how investments in new ideas and innovation drive robust economic growth. Earlier economists had noted the connection between innovation and growth, but Romer provided an exquisite explanation for how it works. In the decades since, Romer's conclusions have been the intellectual inspiration for many in Silicon Valley and help account for how it has attained such wealth.

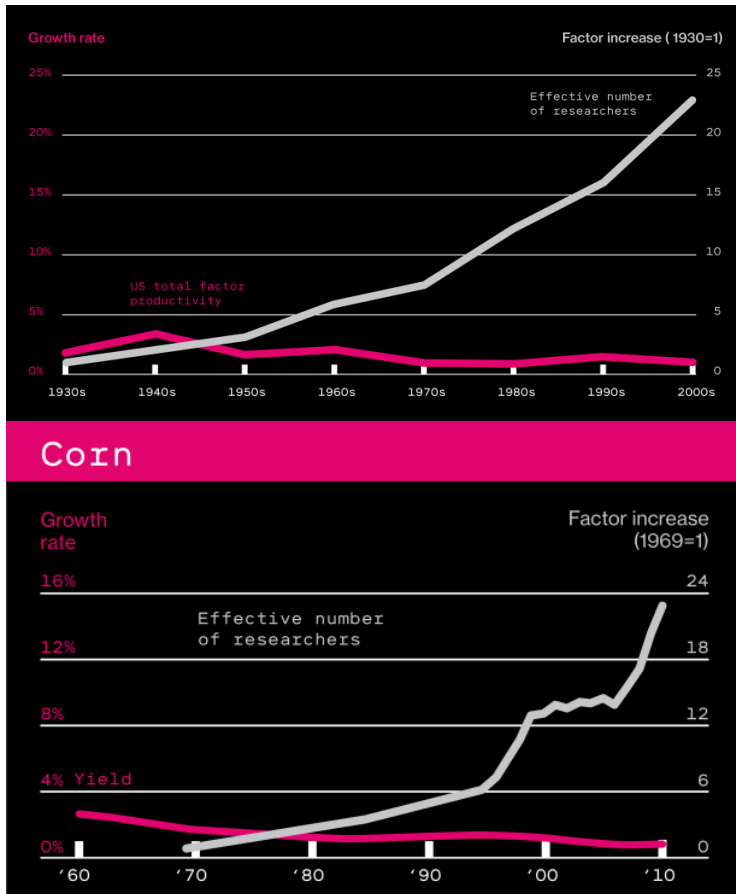
But what if our pipeline of new ideas is drying up? Economists Nicholas Bloom and Chad Jones at Stanford, Michael Webb, a graduate student at the university, and John Van Reenen at MIT looked at the problem in a recent paper called "Are ideas getting harder to find?" (Their answer was "Yes.") Looking at drug discovery, semiconductor research, medical innovation, and efforts to improve crop yields, the economists found a common story: investments in research are climbing sharply, but the payoffs are staying constant.

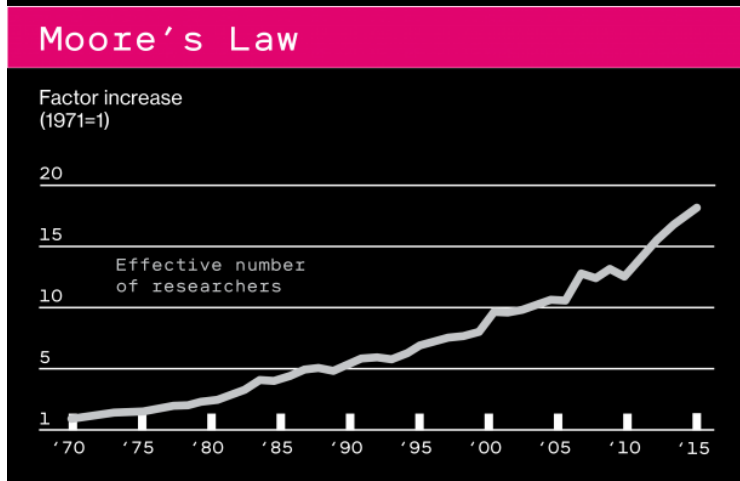
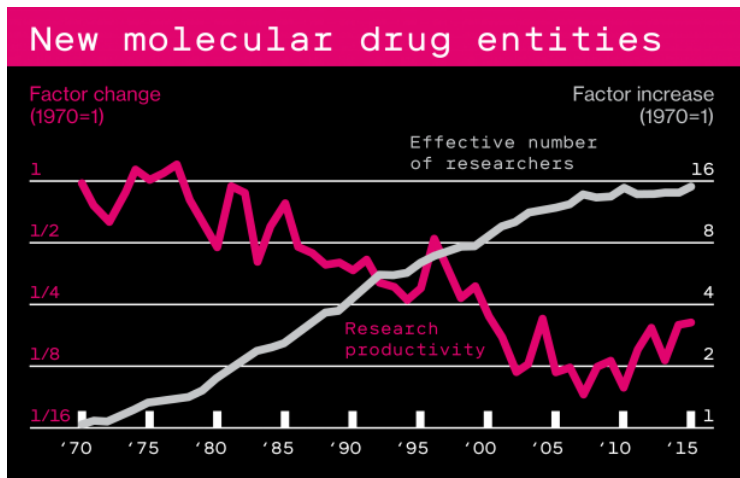
From an economist's perspective, that's a productivity problem: we're paying more for a similar amount of output. And the numbers look bad. Research productivity—the number of researchers it takes to produce a given result—is declining by around 6.8%

annually for the task of extending Moore’s Law, which requires that we find ways to pack ever more and smaller components on a semiconductor chip in order to keep making computers faster and more powerful. (It takes more than 18 times as many researchers to double chip density today as it did in the early 1970s, they found.) For improving seeds, as measured by crop yields, research productivity is dropping by around 5% each year. For the US economy as a whole, it is declining by 5.3%.

The rising price of big ideas

It is taking more researchers and money to find productive new ideas, according to economists at Stanford and MIT. That’s a likely factor in the overall sluggish growth in the US and Europe in recent decades. The graph below shows the pattern for the overall economy, highlighting US total factor productivity (by decade average and for 2000–2014)—a measure of the contribution of innovation—versus the number of researchers. Similar patterns hold for specific research areas.





SOURCE: BLOOM, JONES, VAN REENEN, AND WEBB

Any negative effect of this decline has been offset, so far, by the fact that we're putting more money and people into research. So we're still doubling the number of transistors on a chip every two years, but only because we're dedicating far more people to the problem. We'll have to double our investments in research and development over the next 13 years just to keep trading water.

It could be, of course, that fields like crop science and semiconductor research are getting old and the opportunities for innovation are shriveling up. However, the researchers also found that overall growth tied to innovation in the economy was slow. Any investments in new areas, and any inventions they have generated, have failed to change the overall story.

The drop in research productivity appears to be a decades-long trend. But it is particularly worrisome to economists now because we've seen an overall slowdown in economic growth since the mid-2000s. At a time of brilliant new technologies like smartphones, driverless cars, and Facebook, growth is sluggish, and the portion of it attributed to innovation—called total factor productivity—has been particularly weak.

The lingering effects of the 2008 financial collapse could be hampering growth, says Van Reenen, and so could continuing political uncertainties. But dismal research productivity is undoubtedly a contributor. And he says that if the decline continues, it could do serious damage to future prosperity and growth.

It makes sense that we've already picked much of what some economists like to call the "low-hanging fruit" in terms of inventions. Could it be that the only fruit left is a few shriveled apples on the farthest branches of the tree? Robert Gordon, an economist at Northwestern University, has been a strong proponent of that view. He says we're unlikely to match the flourishing of discovery that marked the late 19th and early 20th centuries, when inventions such as electric light and power and the internal-combustion engine led to a century of unprecedented prosperity.

If Gordon is right, and there are fewer big inventions left, we're doomed to a dismal economic future. But few economists think that's the case. Rather, it makes sense that big new ideas are out there; it's just getting more expensive to find them as the science becomes increasingly complex. The chances that the next penicillin will just fall into our laps are slim. We'll need more and more researchers to make sense of the advancing science in fields like chemistry and biology.

It's what Ben Jones, an economist at Northwestern, calls "the burden of knowledge." Researchers are becoming more specialized, making it necessary to form larger—and more expensive—teams to solve problems. Jones's research shows that the age at which scientists reach their peak productivity is going up: it takes them longer to gain the expertise they need. "It's an innate by-product of the exponential growth of knowledge," he says.

"A lot of people tell me our findings are depressing, but I don't see it that way," says Van Reenen. Innovation might be more difficult and expensive, but that, he says, simply points to the need for policies, including tax incentives, that will encourage investments into more research.

"As long as you put resources into R&D, you can maintain healthy productivity growth," says Van Reenen. "But we have to be prepared to spend money to do it. It doesn't come free."

Giving up on science

Can AI creatively solve the kinds of problems that such innovation requires? Some experts are now convinced that it can, given the kinds of advances shown off by the game-playing machine AlphaGo.

AlphaGo mastered the ancient game of Go, beating the reigning champion, by studying the nearly unlimited possible moves in a game that has been played for several thousand years by humans relying heavily on intuition. In doing so, it sometimes came up with winning strategies that no human player had thought to try. Likewise, goes the thinking, deep-learning programs trained on large amounts of experimental data and chemical literature could come up with novel compounds that scientists never imagined.

Might an AlphaGo-like breakthrough help the growing armies of researchers poring over ever-expanding scientific data? Could AI make basic research faster and more productive, reviving areas that have become too expensive for businesses to pursue?

The last several decades have seen a massive upheaval in our R&D efforts. Since the days when AT&T's Bell Labs and Xerox's PARC produced world-changing inventions like the transistor, solar cells, and laser printing, most large companies in the US and other rich economies have given up on basic research. Meanwhile, US federal R&D investments have been flat, particularly for fields other than life sciences. So while we continue to increase the number of researchers overall and to turn incremental advances into commercial opportunities, areas that require long-term research and a grounding in basic science have taken a hit.

The invention of new materials in particular has become a commercial backwater. That has held back needed innovations in clean tech—stuff like better batteries, more efficient solar cells, and catalysts to make fuels directly from sunlight and carbon dioxide (think artificial photosynthesis). While the prices of solar panels and batteries are falling steadily, that's largely because of improvements in manufacturing and economies of scale, rather than fundamental advances in the technologies themselves.

Might an AlphaGo-like breakthrough help the growing armies of researchers poring over ever-expanding scientific data?

It takes an average of 15 to 20 years to come up with a new material, says Tonio Buonassisi, a mechanical engineer at MIT who is working with a team of scientists in

Singapore to speed up the process. That’s far too long for most businesses. It’s impractical even for many academic groups. Who wants to spend years on a material that may or may not work? This is why venture-backed startups, which have generated much of the innovation in software and even biotech, have long given up on clean tech: venture capitalists generally need a return within seven years or sooner.

“A 10x acceleration [in the speed of materials discovery] is not only possible, it is necessary,” says Buonassisi, who runs a photovoltaic research lab at MIT. His goal, and that of a loosely connected network of fellow scientists, is to use AI and machine learning to get that 15-to-20-year time frame down to around two to five years by attacking the various bottlenecks in the lab, automating as much of the process as possible. A faster process gives the scientists far more potential solutions to test, allows them to find dead ends in hours rather than months, and helps optimize the materials. “It transforms how we think as researchers,” he says.

It could also make materials discovery a viable business pursuit once again. Buonassisi points to a chart showing the time it took to develop various technologies. One of the columns labeled “lithium-ion batteries” shows 20 years.

Another, much shorter column is labeled “novel solar cell”; at the top is “2030 climate target.” The point is clear: we can’t wait another 20 years for the next breakthrough in clean-tech materials.

AI startups in drugs and materials

| | 1 Atomwise | 2 Kebotix | 3 Deep Genomics |
|-----------------------|---|--|---|
| What they do | Use neural networks to search through large databases to find small drug-like molecules that bind to targeted proteins. | Develop a combination of robotics and AI to speed up the discovery and development of new materials and chemicals. | Use artificial intelligence to search for oligonucleotide molecules to treat genetic diseases. |
| Why it matters | Identifying such molecules with desirable properties, such as potency, is a critical first step in drug discovery. | It takes more than a decade to develop a material. Cutting that time could help us tackle problems such as climate change. | Oligonucleotide treatments hold promise against a range of diseases, including neurodegenerative and metabolic disorders. |

The AI-driven lab

“Come to a free land”: that is how Alán Aspuru-Guzik invites a US visitor to his Toronto lab these days. In 2018 Aspuru-Guzik left his tenured position as a Harvard chemistry professor, moving with his family to Canada. His decision was driven by a strong distaste for President Donald Trump and his policies, particularly those on immigration. It didn’t hurt, however, that Toronto is rapidly becoming a mecca for artificial-intelligence research.

As well as being a chemistry professor at the University of Toronto, Aspuru-Guzik also has a position at the Vector Institute for Artificial Intelligence. It’s the AI center cofounded by Geoffrey Hinton, whose pioneering work on deep learning and neural networks is largely credited with jump-starting today’s boom in AI.

In a notable 2012 paper, Hinton and his coauthors demonstrated that a deep neural network, trained on a huge number of pictures, could identify a mushroom, a leopard, and a dalmatian dog. It was a remarkable breakthrough at the time, and it quickly ushered in an AI revolution using deep-learning algorithms to make sense of large data sets.

Researchers rapidly found ways to use such neural networks to help driverless cars navigate and to spot faces in a crowd. Others modified the deep-learning tools so that they could train themselves; among these tools are GANs (generative adversarial networks), which can fabricate images of scenes and people that never existed.

In a 2015 follow-up paper, Hinton provided clues that deep learning could be used in chemistry and materials research. His paper touted the ability of neural network to discover “intricate structures in high-dimensional data”—in other words, the same networks that can navigate through millions of images to find, say, a dog with spots could sort through millions of molecules to identify one with certain desirable properties.

Energetic and bubbling with ideas, Aspuru-Guzik is not the type of scientist to patiently spend two decades figuring out whether a material will work. And he has quickly adapted deep learning and neural networks to attempt to reinvent materials discovery. The idea is to infuse artificial intelligence and automation into all the steps of materials research: the initial design and synthesis of a material, its testing and analysis, and finally the multiple refinements that optimize its performance.

On a freezing cold day early this January, Aspuru-Guzik has his hat pulled tightly down over his ears but otherwise seems oblivious to the bitter Canadian weather. He has other things on his mind. For one thing, he's still waiting for the delivery of a \$1.2 million robot, now on a ship from Switzerland, that will be the centerpiece for the automated, AI-driven lab he has envisioned.

In the lab, deep-learning tools like GANs and their cousin, a technique called autoencoder, will imagine promising new materials and figure out how to make them. The robot will then make the compounds; Aspuru-Guzik wants to create an affordable automated system that would be able to spit out new molecules on demand. Once the materials have been made, they can be analyzed with instruments such as a mass spectrometer. Additional machine-learning tools will make sense of that data and “diagnose” the material's properties. These insights will then be used to further optimize the materials, tweaking their structures. And then, Aspuru-Guzik says, “AI will select the next experiment to make, closing the loop.”

The idea is to infuse artificial intelligence and automation into all the steps of materials research and drug discovery.

Once the robot is in place, Aspuru-Guzik expects to make some 48 novel materials every two days, drawing on the machine-learning insights to keep improving their structures. That's one promising new material every hour, an unprecedented pace that could completely transform the lab's productivity.

It's not all about simply dreaming up “a magical material,” he says. To really change materials research, you need to attack the entire process: “What are the bottlenecks? You want AI in every piece of the lab.” Once you have a proposed structure, for example, you still need to figure out how to make it. It can take weeks to months to solve what chemists call “retrosynthesis”—working backwards from a molecular structure to figure out the steps needed to synthesize such a compound. Another bottleneck comes in making sense of the reams of data produced by analytic equipment. Machine learning could speed up each of those steps.

What motivates Aspuru-Guzik is the threat of climate change, the need for improvements in clean tech, and the essential role of materials in producing such advances. His own research is looking at novel organic electrolytes for flow batteries, which can be used to store excess electricity from power grids and pump it back in

when it's needed, and at organic solar cells that would be far cheaper than silicon-based ones. But if his design for a self-contained, automated chemical lab works, he suggests, it could make chemistry far more accessible to almost anyone. He calls it the "democratization of materials discovery."

"This is where the action is," he says. "AIs that drive cars, AIs that improve medical diagnostics, AIs for personal shopping—the economic growth from AIs applied to scientific research may swamp the economic impact from all those other AIs combined."

The Vector Institute, Toronto's magnet for AI research, sits less than a mile away. From the windows of the large open office space, you can look across at Ontario's parliament building. The proximity of experts in AI, chemistry, and business to the province's seat of government in downtown Toronto isn't a coincidence. There's a strong belief among many in the city that AI will transform business and the economy, and increasingly, some are convinced it will radically change how we do science.

Still, if it is do that, a first step is convincing scientists it is worthwhile.

Amgen's Guzman-Perez says many of his peers in medicinal chemistry are skeptical. Over the last few decades the field has seen a series of supposedly revolutionary technologies, from computational design to combinatorial chemistry and high-throughput screening, that have automated the rapid production and testing of multiple molecules. Each has proved somewhat helpful but limited. None, he says, "magically get you a new drug."

It's too early to know for sure whether deep learning could finally be the game-changer, he acknowledges, "and it's hard to know the time frame." But he takes encouragement from the speed at which AI has transformed image recognition and other search tasks.

https://www.technologyreview.com/s/612898/ai-is-reinventing-the-way-we-invent/?utm_medium=tr_social&utm_campaign=site_visitor.unpaid.engagement&utm_source=LinkedIn#Echobox=1573150688