The Simple Economics of Machine Intelligence

By Ajay Agrawal, Joshua Gans and Avi Goldfarb

November 17, 2016

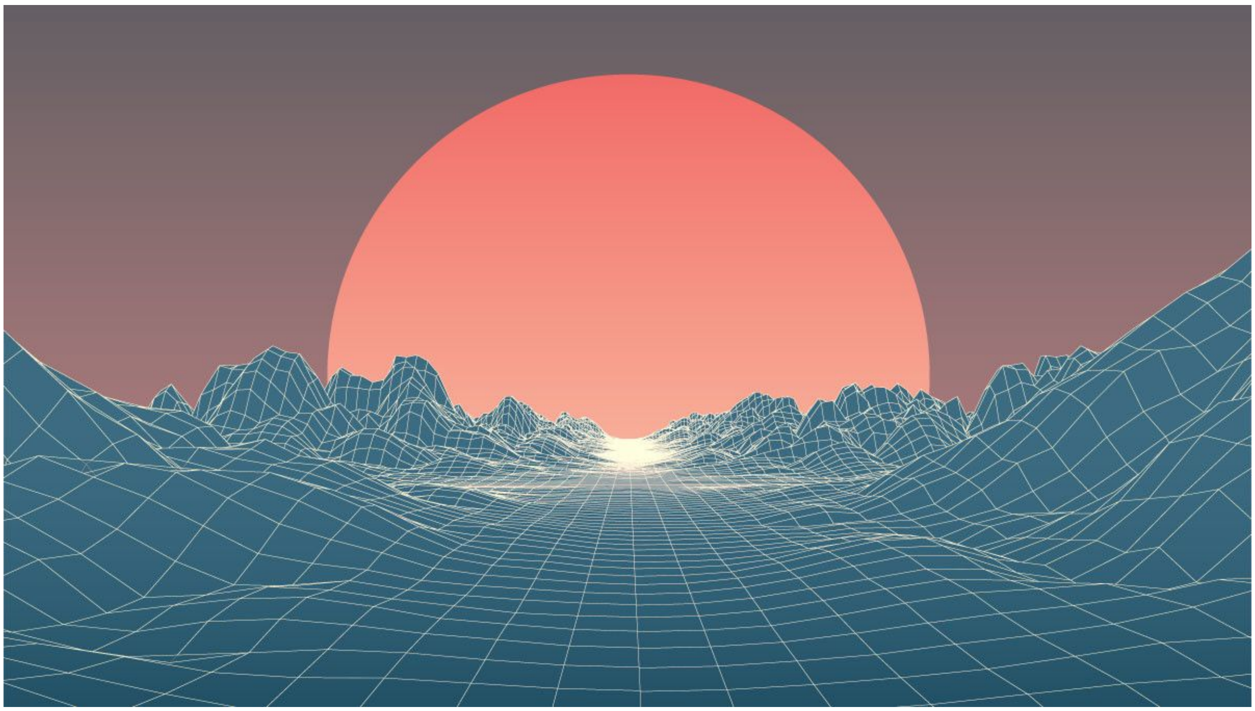

The year 1995 was heralded as the beginning of the "New Economy." Digital communication was set to upend markets and *change everything*. But economists by and large didn't buy into the hype. It wasn't that we didn't recognize that something changed. It was that we recognized that the old economics lens remained useful for looking at the changes taking place. The economics of the "New Economy" could be described at a high level: Digital technology would cause a reduction in the cost of search and communication. This would lead to more search, more communication, and more activities that go together with search and communication. That's essentially what happened.

Today we are seeing similar hype about machine intelligence. But once again, as economists, we believe some simple rules apply. Technological revolutions tend to involve some important activity becoming cheap, like the cost of communication or finding information. Machine intelligence is, in its essence, a prediction technology, so the economic shift will center around a drop in the cost of prediction.

The first effect of machine intelligence will be to lower the cost of goods and services that rely on prediction. This matters because prediction is an input to a host of activities including transportation, agriculture, healthcare, energy manufacturing, and retail.

When the cost of any input falls so precipitously, there are two other well-established economic implications. First, we will start using prediction to perform tasks where we previously didn't. Second, the value of other things that complement prediction will rise.

Lots of tasks will be reframed as prediction problems

As machine intelligence lowers the cost of prediction, we will begin to use it as an input for things for which we never previously did. As a historical example, consider semiconductors, an area of technological advance that caused a significant drop in the cost of a different input: arithmetic. With semiconductors we could calculate cheaply, so activities for which arithmetic was a key input, such as data analysis and accounting, became much cheaper. However, we also started using the newly cheap arithmetic to solve problems that were not historically arithmetic problems. An example is photography. We shifted from a film-oriented, chemistry-based approach to a digital-oriented, arithmetic-based approach. Other new applications for cheap arithmetic include communications, music, and drug discovery.

The same goes for machine intelligence and prediction. As the cost of prediction falls, not only will activities that were historically prediction-oriented become cheaper — like inventory management and demand forecasting — but we will also use prediction to tackle other problems for which prediction was not historically an input.

Consider navigation. Until recently, autonomous driving was limited to highly controlled environments such as warehouses and factories where programmers could anticipate the range of scenarios a vehicle may encounter, and could program if-then-else-type decision algorithms accordingly (e.g., "If an object approaches the vehicle, then slowdown"). It was inconceivable to put an autonomous vehicle on a city street because the number of possible scenarios in such an uncontrolled environment would require programming an almost infinite number of if-then-else statements.

Inconceivable, that is, until recently. Once prediction became cheap, innovators reframed driving as a prediction problem. Rather than programing endless if-then-else statements, they instead simply asked the AI to predict: "What would a human driver do?” They outfitted vehicles with a variety of sensors – cameras, lidar, radar, etc. – and then collected millions of miles of human driving data. By linking the incoming environmental data from sensors on the outside of the car to the driving decisions made by the human inside the car (steering, braking, accelerating), the AI learned to predict how humans would react to each second of incoming data about their environment. Thus, prediction is now a major component of the solution to a problem that was previously not considered a prediction problem.

Judgment will become more valuable

When the cost of a foundational input plummets, it often affects the value of other inputs. The value goes up for complements and down for substitutes. In the case of photography, the value of the hardware and software components associated with digital cameras went up as the cost of arithmetic dropped because demand increased – we wanted more of them. These components were complements to arithmetic; they were used together. In contrast, the value of film-related chemicals fell – we wanted less of them.

All human activities can be described by five high-level components: data, prediction, judgment, action, and outcomes. For example, a visit to the doctor in response to pain leads to: 1) x-rays, blood tests, monitoring (data), 2) diagnosis of the problem, such as “if we administer treatment A, then we predict outcome X, but if we administer treatment B, then we predict outcome Y” (prediction), 3) weighing options: “given your age, lifestyle, and family status, I think you might be best with treatment A; let’s discuss how you feel about the risks and side effects” (judgment); 4) administering treatment A (action), and 5) full recovery with minor side effects (outcome).

As machine intelligence improves, the value of human prediction skills will decrease because machine prediction will provide a cheaper and better substitute for human prediction, just as machines did for arithmetic. However, this does not spell doom for human jobs, as many experts suggest. That’s because the value of human judgment skills will increase. Using the language of economics, judgment is a complement to prediction and therefore when the cost of prediction falls demand for judgment rises. We’ll want more human judgment.

For example, when prediction is cheap, diagnosis will be more frequent and convenient, and thus we’ll detect many more early-stage, treatable conditions. This will mean more decisions will be made about medical treatment, which means greater demand for the application of ethics, and for emotional support, which are provided by humans. The line between judgment and prediction isn’t clear cut – some judgment tasks will even be reframed as a series of predictions. Yet, overall the value of prediction-related human skills will fall, and the value of judgment-related skills will rise.

Interpreting the rise of machine intelligence as a drop in the cost of prediction doesn't offer an answer to every specific question of how the technology will play out. But it yields two key implications: 1) an expanded role of prediction as an input to more goods and services, and 2) a change in the value of other inputs, driven by the extent to which they are complements to or substitutes for prediction. These changes are coming. The speed and extent to which managers should invest in judgment-related capabilities will depend on the how fast the changes arrive.

How AI Will Change Strategy: A Thought Experiment

By Ajay Agrawal, Joshua Gans and Avi Goldfarb

October 3, 2019

How will AI change strategy? That's the single most common question the three of us are asked from corporate executives, and it's not trivial to answer. AI is fundamentally a prediction technology. As advances in AI make prediction cheaper, economic theory dictates that we'll use prediction more frequently and widely, and the value of complements to prediction – like human judgment – will rise. But what does all this mean for strategy?

Here's a thought experiment we've been using to answer that question. Most people are familiar with shopping at Amazon. Like with most online retailers, you visit their website, shop for items, place them in your "basket," pay for them, and then Amazon ships them to you. Right now, Amazon's business model is shopping-then-shipping.

Most shoppers have noticed Amazon's recommendation engine while they shop — it offers suggestions of items that their AI predicts you will want to buy. At present, Amazon's AI does a reasonable job, considering the millions of items on offer. However, they are far from perfect. In our case, the AI accurately predicts what we want to buy about 5% of the time. In other words, we actually purchase about one out of every 20 items it recommends. Not bad!

Now for the thought experiment. Imagine the Amazon AI collects more information about us: in addition to our searching and purchasing behavior on their website, it also collects other data it finds online, including social media, as well as offline, such as our shopping behavior at Whole Foods. It knows not only what we buy, but also what time we go to the store, which location we shop at, how we pay, and more.

Now, imagine the AI uses that data to improve its predictions. We think of this sort of improvement as akin to turning up the volume knob on a speaker dial. But rather than volume, you're turning up the AI's prediction accuracy. What happens to Amazon's strategy as their data scientists, engineers, and machine learning experts work tirelessly to dial up the accuracy on the prediction machine?

At some point, as they turn the knob, the AI's prediction accuracy crosses a threshold, such that it becomes in Amazon's interest to change its business model. The prediction becomes sufficiently accurate that it becomes more profitable for Amazon to ship you the goods that it predicts you will want rather than wait for you to order them. Every week, Amazon ships you boxes of items it predicts you will want, and then you shop in the comfort and convenience of your own home by choosing the items you wish to keep from the boxes they delivered.

This approach offers two benefits to Amazon. First, the convenience of predictive shipping makes it much less likely that you purchase the items from a competing retailer as the products are conveniently delivered to your home before you buy them elsewhere. Second, predictive shipping nudges you to buy items that you were considering purchasing but might not have gotten around to. In both cases, Amazon gains a higher share-of-wallet. Turning the prediction dial up far enough changes Amazon's business model from shopping-then-shipping to shipping-then-shopping.

Of course, shoppers would not want to deal with the hassle of returning all the items they don't want. So, Amazon would invest in infrastructure for the product returns —

perhaps a fleet of delivery-style trucks that do pick-ups once a week, conveniently collecting items that customers don't want.

If this is a better business model, then why hasn't Amazon done it already? Well, they may be working on it. But if it were implemented today, the cost of collecting and handling returned items would outweigh the increase in revenue from a greater share-of-wallet. For example, today we would return 95% of the items it ships to us. That is annoying for us and costly for Amazon. The prediction isn't good enough for Amazon to adopt the new model.

That said, one can imagine a scenario where Amazon adopts the new strategy even *before* the prediction accuracy is good enough to make it profitable because the company *anticipates* that at some point it will be profitable. By launching sooner, Amazon's AI will get more data sooner, and improve faster. Amazon realizes that the sooner it gets started, the harder it will be for competitors to catch up. Better predictions will attract more shoppers, more shoppers will generate more data to train the AI, more data will lead to better predictions, and so on, creating a virtuous circle. In other words, there are increasing returns to AI, and thus the timing of adopting this kind of strategy matters. Adopting too early could be costly, but adopting too late could be fatal.

The key insight here is that turning the dial on the prediction machine has a significant impact on strategy. In this example, it shifts Amazon's business model from shopping-then-shipping to shipping-then-shopping, generates the incentive to vertically integrate into operating a product-returns service (including a fleet of trucks), and accelerates the timing of investment due to first-mover advantage from increasing returns. All this is due to the single act of turning the dial on the prediction machine.

Most readers will be familiar with the outcome of companies like Blockbuster and Borders that underestimated how quickly the online consumer behavior dial would turn in the context of online shopping and the digital distribution of goods and services. Perhaps they were lulled into complacency by the initially slow adoption rate of this technology in the early days of the commercial internet (1995-1998).

Today, in the case of AI, some companies are making early bets anticipating that the dial on the prediction machine will start turning faster once it gains momentum. Most people are familiar with Google's 2014 acquisition of DeepMind – over $500M for a company that had generated negligible revenue, but had developed an AI that learned to play certain Atari games at a super human performance level. Perhaps fewer readers are aware that more traditional companies are also making bets on the pace the dial will turn. In 2016, GM paid over $1B to acquire AI startup Cruise Automation, and in 2017, Ford invested $1B in AI startup Argo AI, and John Deere paid over $300M to acquire AI startup Blue River Technology – all three startups had generated negligible revenue relative to the price at the time of purchase. GM, Ford, and John Deere are each betting on an exponential speed up of AI performance and, at those prices, anticipating a significant impact on their business strategies.

Strategists face two questions in light of all of this. First, they must invest in developing a better understanding of how fast and how far the dial on their prediction machines will turn for their sector and applications. Second, they must invest in developing a thesis about the strategy options created by the shifting economics of their business that result from turning the dial, similar to the thought experiment we considered for Amazon.

So, the overarching theme for initiating an AI strategy? Close your eyes, imagine putting your fingers on the dial of your prediction machine, and, in the immortal words of Spinal Tap, turn it to eleven.

The ideas here are adapted from our forthcoming book "Prediction Machines: The Simple Economics of Artificial Intelligence." (Harvard Business School Press, April 2018)

The Trade-Off Every AI Company Will Face

By Ajay Agrawal, Joshua Gans and Avi Goldfarb

March 28, 2017

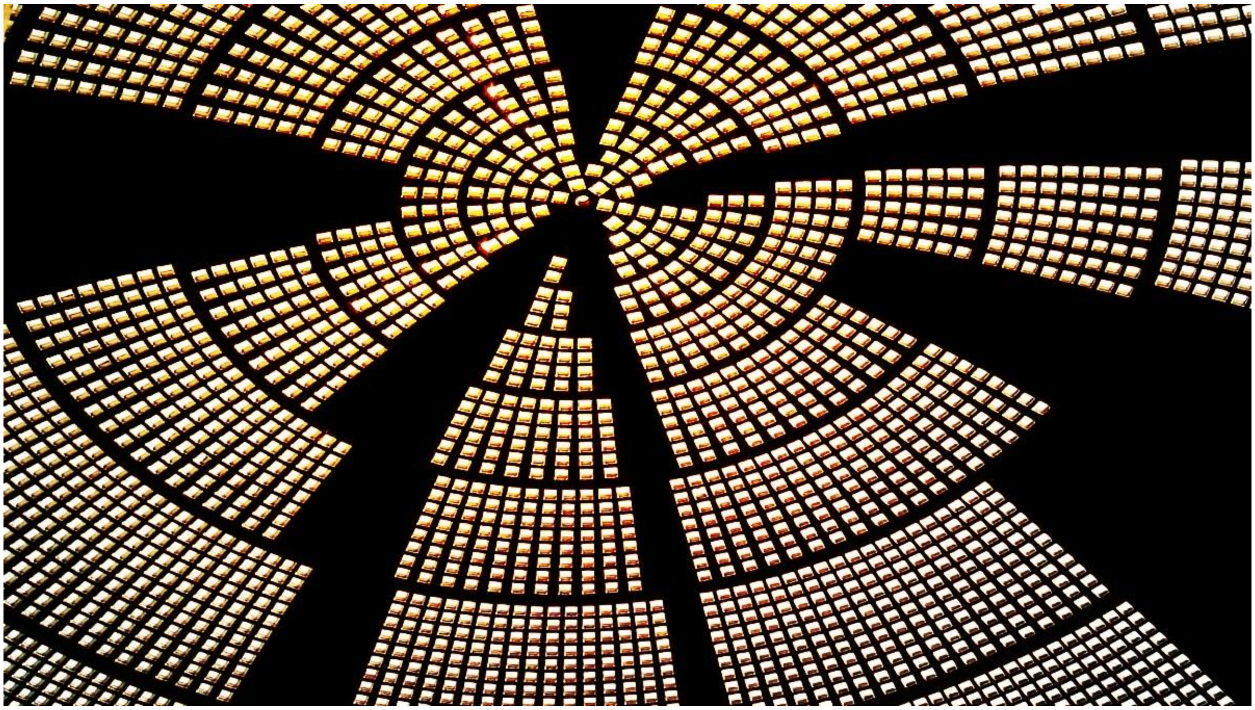

It doesn't take a tremendous amount of training to begin a job as a cashier at McDonald's. Even on their first day, most new cashiers are good enough. And they improve as they serve more customers. Although a new cashier may be slower and make more mistakes than their experienced peers, society generally accepts that they will learn from experience.

We don't often think of it, but the same is true of commercial airline pilots. We take comfort that airline transport pilot certification is regulated by the U.S. Department of Transportation's Federal Aviation Administration and requires minimum experience of

1,500 hours of flight time, 500 hours of cross-country flight time, 100 hours of night flight time, and 75 hours of instrument operations time. But we also know that pilots continue to improve from on-the-job experience.

On January 15, 2009, when US Airways Flight 1549 was struck by a flock of Canada geese, shutting down all engine power, Captain Chelsey “Sully” Sullenberger miraculously landed his plane in the Hudson River, saving the lives of all 155 passengers. Most reporters attributed his performance to experience. He had recorded 19,663 total flight hours, including 4,765 flying an A320. Sully himself reflected: “One way of looking at this might be that for 42 years, I’ve been making small, regular deposits in this bank of experience, education, and training. And on January 15, the balance was sufficient so that I could make a very large withdrawal.” Sully, and all his passengers, benefited from the thousands of people he’d flown before.

The difference between cashiers and pilots in what constitutes “good enough” is based on tolerance for error. Obviously, our tolerance is much lower for pilots. This is reflected in the amount of in-house training we require them to accumulate prior to serving their first customers, even though they continue to learn from on-the-job experience. We have different definitions for good enough when it comes to how much training humans require in different jobs.

The same is true of machines that learn.

Artificial intelligence (AI) applications are based on generating predictions. Unlike traditionally programmed computer algorithms, designed to take data and follow a specified path to produce an outcome, machine learning, the most common approach to AI these days, involves algorithms evolving through various learning processes. A machine is given data, including outcomes, it finds associations, and then, based on those associations, it takes new data it has never seen before and predicts an outcome.

This means that intelligent machines need to be trained, just as pilots and cashiers do. Companies design systems to train new employees until they are good enough and then deploy them into service, knowing that they will improve as they learn from experience doing their job. While this seems obvious, determining what constitutes good enough is an important decision. In the case of machine intelligence, it can be a major strategic decision regarding timing: when to shift from in-house training to on-the-job learning.

There is no ready-made answer as to what constitutes “good enough” for machine intelligence. Instead, there are trade-offs. Success with machine intelligence will require taking these trade-offs seriously and approaching them strategically.

The first question firms must ask is what tolerance they and their customers have for error. We have high tolerance for error with some intelligent machines and a low tolerance for others. For example, Google's Inbox application reads your email, uses AI to predict how you will want to respond, and generates three short responses for the user to choose from. Many users report enjoying using the application even when it has a 70% failure rate (i.e., the AI-generated response is only useful 30% of the time). The reason for this high tolerance for error is that the benefit of reduced composing and typing outweighs the cost of wasted screen real estate when the predicted short response is wrong.

In contrast, we have low tolerance for error in the realm of autonomous driving. The first generation of autonomous vehicles, largely pioneered by Google, was trained using specialist human drivers who took a limited set of vehicles and drove them hundreds of thousands of kilometers. It was like a parent taking a teenager on supervised driving experiences before letting them drive on their own.

The human specialist drivers provide a safe training environment, but are also extremely limited. The machine only learns about a small number of situations. It may take many millions of miles in varying environments and situations before someone has learned how to deal with the rare incidents that are more likely to lead to accidents. For autonomous vehicles, real roads are nasty and unforgiving precisely because nasty or unforgiving human-caused situations can occur on them.

The second question to ask, then, is how important it is to capture user data in the wild. Understanding that training might take a prohibitively long time, Tesla rolled out autonomous vehicle capabilities to all its recent models. These capabilities included a set of sensors that collect environmental data as well as driving data that is uploaded to Tesla's machine learning servers. In a very short period of time, Tesla can obtain training data just by observing how the drivers of its cars drive. The more Tesla vehicles there are on the roads, the more Tesla's machines can learn.

However, in addition to passively collecting data as humans drive their Teslas, the company needs autonomous driving data to understand how its autonomous systems are operating. For that, it needs to have cars drive autonomously so that it can assess performance, but also assess when a human driver, required to be there and paying attention, chooses to intervene. Tesla's ultimate goal is not to produce a copilot, or a teenager who drives under supervision, but a fully autonomous vehicle. That requires getting to the point where real people feel comfortable in a self-driving car.

Herein lies a tricky trade-off. In order to get better, Tesla needs its machines to learn in real situations. But putting its current cars in real situations means giving customers a relatively "young and inexperienced" driver — although perhaps as good as or better than many young human drivers. Still, this is far riskier than beta testing, for example, whether Siri or Alexa understood what you said, or whether Google Inbox correctly predicts your response to an email. In the case of Siri, Alexa, or Google Inbox, it means a lower-quality user experience. In the case of autonomous vehicles, it means putting lives at risk.

As Backchannel documented in a recent article, that experience can be scary. Cars can exit freeways without notice, or put on the brakes when mistaking an underpass for an obstruction. Nervous drivers may opt not to use the autonomous features, and, in the process, may hinder Tesla's ability to learn. Furthermore, even if the company can persuade some people to become beta testers, are those the people it wants? After all, a beta tester for autonomous driving may be someone with a taste for more risk than the average driver. In that case, who is the company training their machines to be like?

Machines learn faster with more data, and more data is generated when machines are deployed in the wild. However, bad things can happen in the wild and harm the company brand. Putting products in the wild earlier accelerates learning but risks harming the brand (and perhaps the customer!); putting products in the wild later slows learning but allows for more time to improve the product in-house and protect the brand (and, again, perhaps the customer).

For some products, like Google Inbox, the answer to the trade-off seems clear because the cost of poor performance is low and the benefits from learning from customer usage are high. It makes sense to deploy this type of product in the wild early. For other products, like cars, the answer is less clear. As more companies seek to take advantage of machine learning, this is a trade-off more and more will have to make.