

Is a Cambrian Explosion Coming for Robotics?

Gill A. Pratt

About half a billion years ago, life on earth experienced a short period of very rapid diversification called the “Cambrian Explosion.” Many theories have been proposed for the cause of the Cambrian Explosion, with one of the most provocative being the evolution of vision, which allowed animals to dramatically increase their ability to hunt and find mates (for discussion, see Parker 2003). Today, technological developments on several fronts are fomenting a similar explosion in the diversification and applicability of robotics. Many of the base hardware technologies on which robots depend—particularly computing, data storage, and communications—have been improving at exponential growth rates. Two newly blossoming technologies—“Cloud Robotics” and “Deep Learning”—could leverage these base technologies in a virtuous cycle of explosive growth. In Cloud Robotics—a term coined by James Kuffner (2010)—every robot learns from the experiences of all robots, which leads to rapid growth of robot competence, particularly as the number of robots grows. Deep Learning algorithms are a method for robots to learn and generalize their associations based on very large (and often cloud-based) “training sets” that typically include millions of examples. Interestingly, Li (2014) noted that one of the robotic capabilities recently enabled by these combined technologies is vision—the same capability that may have played a leading role in the Cambrian Explosion.

■ *Gill A. Pratt will be stepping down next month from his position as a Program Manager of the Defense Advanced Research Projects Agency (DARPA), Arlington, Virginia, where he oversaw the DARPA Robotics Challenge and several other programs in robotics while on leave from the faculty of Franklin W. Olin College, Needham, Massachusetts. The views expressed in this paper are those of the author, not Olin College, DARPA, the US Department of Defense, or the US Government.*

How soon might a Cambrian Explosion of robotics occur? It is hard to tell. Some say we should consider the history of computer chess, where brute force search and heuristic algorithms can now beat the best human player yet no chess-playing program inherently knows how to handle even a simple adjacent problem, like how to win at a straightforward game like tic-tac-toe (Brooks 2015). In this view, specialized robots will improve at performing well-defined tasks, but in the real world, there are far more problems yet to be solved than ways presently known to solve them.

But unlike computer chess programs, where the rules of chess are built in, today's Deep Learning algorithms use general learning techniques with little domain-specific structure. They have been applied to a range of perception problems, like speech recognition and now vision. It is reasonable to assume that robots will in the not-too-distant future be able perform any associative memory problem at human levels, even those with high-dimensional inputs, with the use of Deep Learning algorithms. Furthermore, unlike computer chess, where improvements have occurred at a gradual and expected rate, the very fast improvement of Deep Learning has been surprising, even to experts in the field. The recent availability of large amounts of training data and computing resources on the cloud has made this possible; the algorithms being used have existed for some time and the learning process has actually become simpler as performance has improved.

While the so-called "neural networks" on which Deep Learning is often implemented differ from what is known about the architecture of the brain in several ways, their distributed "connectionist" approach is more similar to the nervous system than previous artificial intelligence techniques (like the search methods used for computer chess). Several characteristics of real brains are yet to be accomplished, such as episodic memory and "unsupervised learning" (the clustering of similar experiences without instruction), but it seems likely that Deep Learning will soon be able to replicate the performance of many of the perceptual parts of the brain. While questions remain as to whether similar methods can also replicate cognitive functions, the architectures of the perceptual and cognitive parts of the brain appear to be anatomically similar. There is thus reason to believe that artificial cognition may someday be put into effect through Deep Learning techniques augmented with short-term memory systems and new methods of doing unsupervised learning. To date, there are no huge datasets about cognition that are similar to the picture and speech datasets that have been so effective for computer learning about perception. But some methodologies for collecting such datasets, described below, may be possible.

The timing of tipping points is hard to predict, and exactly when an explosion in robotics capabilities will occur is not clear. Commercial investment in autonomy and robotics—including and especially in autonomous cars—has significantly accelerated, with high-profile firms like Amazon, Apple, Google, and Uber, as well as all the automotive companies, announcing significant projects in this area. In the next two sections of this paper, I examine some key technologies contributing to the present excitement in the robotics field. As with other technological developments, there has been a significant uptick in concerns about the societal implication of robotics and

artificial intelligence. Thus, I offer some thoughts about how robotics may affect the economy and some ways to address potential difficulties.

Eight Technical Drivers

A number of technologies relevant to the development of robotics are improving at exponential rates. Here, I discuss eight of the most important. The first three technological developments relate to individual robots; the next two relate to connectivity; and the final three relate to the capacities of the Internet that will shape the future of cloud robotics.

1) *Exponential growth in computing performance.* Robots are made up of computers that allow sensors and actuators to collaborate, and the processing power of computers keeps rising. Moore's law, as originally proposed by Gordon Moore (1965, 1975), one of the founders of Intel, refers to the doubling of transistor count on integrated circuits roughly every 18–24 months with similar improvements in processing speed. This relationship has held for many decades, although it is now approaching some fundamental limits. Semiconductor companies are now etching transistors on to chips at a scale of 14 nanometers (for example, see <http://www.intel.com/content/www/us/en/silicon-innovations/intel-14nm-technology.html>), where a nanometer is an almost unimaginably small one-billionth of a meter: for scale, a typical sheet of paper is about 100,000 nanometers thick. This small scale is approaching physical limits because it involves working at close to the atomic level. However, there do appear to be technologies that can continue Moore's law for at least another few years (Bauer, Veira, and Weig 2013), such as those that move beyond two-dimensional integrated circuits and work at a system level with three-dimensional and multi-chip systems.

2) *Improvements in electromechanical design tools and numerically controlled manufacturing tools.* Modern computer-aided design tools have significantly improved the productivity of electromechanical designers, including the quality of what is designed and the sophistication of what can be designed. Numerically controlled manufacturing tools—including the new “additive processes” such as 3D printing—can build such designs with great precision and little cost to additional complexity. Embedded processors have allowed for tremendous control complexity and performance and reliability improvements, all of which increase the competence and reliability of robots.

3) *Improvements in electrical energy storage.* If robots are to be mobile, they need to find ways to store or generate sufficient power to operate for reasonable periods—at least between episodes of recharging. Over the last few decades, advances in electric batteries and fuel cells have had a poor history of living up to promises. Standard lithium-ion batteries still remain nearly an order of magnitude less energy-dense than hydrocarbon fuels (including gasoline, as well as sugars and fats), but the gap is slowing closing. The high demand and fierce competition in the portable electronic markets of laptops, tablets, and cell phones, not to mention hybrid and electric vehicles, has continued to spark innovation and steady improvement in

energy storage. Supercapacitors, a new technology, charge and discharge much faster than standard batteries, and can be recharged hundreds or thousands of times, but significant improvements still need to be made in how much energy they can store. (For an accessible discussion of the tradeoffs between lithium-ion batteries and supercapacitors, see Miret 2013.) For many robot applications with ready access to recharging infrastructure, battery energy storage is already sufficient, and energy storage should continue to improve over time.

4) *Improvements in electronics power efficiency.* Robots running on electric batteries use electronics for power management of motors, and robots with many motors are particularly sensitive to the cost and performance of these electronics. Power-related semiconductors have taken advantage of general technology improvements in the integrated circuit industry and have also become much less expensive due to the continual improvements in portable devices, which all include batteries that are sensitive to the efficiency of power electronics. LED lighting is another rapidly growing market in which power-electronics semiconductors enable the production of light more efficiently with less power. New types of compound semiconductors (gallium nitride and silicon carbide) promise to usher in even higher performance and lower prices. The computational needs for the combination of cloud robotics and Deep Learning are currently being provided by graphics processing units—extremely high-performance computer chips originally developed for video games. In the future, computation may be provided by neuromorphic (brain-inspired) hardware, which often consumes less power.

5) *Exponential expansion of the availability and performance of local wireless digital communications.* Early robots were essentially stand-alone machines. Their capacities to remember and to solve problems were limited by the programming that they could carry around with them. Updating their information or reprogramming them was a costly and time-consuming process. However, flexible web-connected robots offer different possibilities for programming, problem-solving, learning, and updating. High-performance wireless digital communications are becoming ubiquitous, as are products that leverage that infrastructure, including a wide variety of tablets along with more specialized products. For example, the “learning thermostat” produced by Nest not only can be adjusted from your phone, but it also remembers and learns when you turn it higher or lower, and starts to make future adjustments automatically. The Google Chromecast device lets you take any content from your computer or mobile device and show it on your television using a wireless connection. Average worldwide WiFi Speeds, which were 10 megabits per second in 2014, will nearly double by 2018. There were 48 million public WiFi hotspots globally in 2014, a number that is expected to increase by a factor of seven by 2018 (Cisco 2015). The latest WiFi standard (802.11ac) surpasses one gigabit per second, as does the latest cellular data standard (5G). With these changes, robots will be able to communicate wirelessly within their facilities more quickly than ever.

6) *Exponential growth in the scale and performance of the Internet.* As wireless communication within facilities evolves, so will Internet communication outside the facility. The global Internet is presently estimated to carry about 88 exabytes (that

is, 10^{18} bytes) of traffic per month, which is predicted to double in three years, with no saturation in sight. There are currently about 13 billion devices connected to the Internet, already two for each human being on Earth; this ratio is projected to reach three for each human being on Earth by 2019 (Cisco 2015).

7) *Exponential growth of worldwide data storage.* On a global basis, total information stored is on the order of a 10^{21} bytes,¹ with volume ever increasing due to explosive demand for entertainment and social media. By comparison, the human brain has on the order of 10^{14} synapses. If we crudely equate each synapse to one byte of storage, the current level of worldwide data storage is on the order of 10 million human brains. Of course, bytes on a spinning disk drive accessed serially by a computer are not densely interconnected the way synapses are inside of a brain, so this remarkable amount of information storage has not been combined to emulate even one brain, much less 10 million. But it is a lot of data.

8) *Exponential growth in global computation power.* Worldwide total computation performance has reached on the order of a 10^{21} instructions per second.² More importantly, many billions of disk drives have been produced (although perhaps only a billion or so are running now), and several large Internet companies run millions of high-performance servers in parallel, each with high-performance multiple core processors. As a result, any computations that can be broken into parallel operations—because no communication of intermediate results are needed to solve the separate problems—can be parceled out and solved quickly. (Problems that can be broken into a number of pieces that can be solved in parallel are sometimes called “embarrassingly parallel” problems.) Many problems in robot autonomy can be solved this way.

Cloud Robotics

These technological developments suggest that the capabilities of robots themselves are increasing rapidly, along with the ability to interconnect robots. Several big ideas, collectively known as “cloud robotics,” are poised to leverage many of these technologies to enable the revolution in robot capabilities. The potential gains from cloud robotics can be summarized with four big ideas.

Big Idea #1: Memory-Based Autonomy

The exponential growth in computing and storage performance has led researchers to explore memory-based methods of solving the perception, planning, and control problems relevant to the development of additional degrees of robot autonomy. Instead of decomposing these tasks into a set of hand-coded algorithms customized

¹ Hilbert and Lopez (2011) estimate total global information stored at 2.9×10^{20} bytes in 2007, and growing at 23 percent per year. Extrapolating this growth rate through 2015 provides a basis for the rough estimate in the text.

² Hilbert and Lopez (2011) estimate world computational power at 6.4×10^{18} instructions per second in 2007, and growing at 58 percent per year. Extrapolating this growth rate through 2015 provides a basis for the rough estimate in the text.

for particular circumstances, large numbers of memories of prior experiences can be searched, and a solution based on matching prior experience is used to guide response. When no matching prior memory exists, actions from similar prior memories can be interpolated, or human help can be requested, with the human-provided answer then recorded for future use both at that robot and at other robots as well.

The continuing progress on fast search algorithms for Internet information has accelerated memory-based approaches. Using cloud-based computing, the lookup of prior examples can exploit a large number of external computing resources in parallel with a relatively small amount of communication. But for memory-based methods to work, where do the memories of solutions come from?

Big Idea #2: High-Speed Sharing of Experiences

A single robot, using a memory-based method to implement additional degrees of autonomy, would, like a newborn child, quite probably take decades to learn to do anything useful. Indeed, the robot would be much slower than a human child, because even instincts would be missing.

But while the communications bandwidth inside of a human brain is high relative to that of a robot, human beings communicate externally with one another relatively slowly, at rates on the order of 10 bits per second. Robots, and computers in general, can communicate at rates over one gigabit per second—or roughly 100 million times faster. Based on this tremendous difference in external communication speeds, a combination of wireless and Internet communication can be exploited to share what is learned by every robot with all robots. Human beings take decades to learn enough to add meaningfully to the compendium of common knowledge. However, robots not only stand on the shoulders of each other’s learning, but can start adding to the compendium of robot knowledge almost immediately after their creation.

It is unclear whether future robots will have high-performance on-board “brains” that cache memories from the cloud, using the Internet only when they aren’t sure what to do, or whether high-speed internet communications will mean that most of a robot’s intelligence (like most of the computation done by some modern video games) will be accomplished remotely using computational resources on the cloud. In either case, the capacities of fast communication and the Internet will catalyze increases in robot capability.

Big Idea #3: Learning from Imagination

Human beings often use imagination to practice and prepare for future circumstances. Similarly, a robot—or a cloud-computing robot “brain”—can use simulation to explore circumstances that may be faced by a robot in the future and to experiment with possible solutions, remembering only those that worked. Such simulations can be done without the need for any physical activity, and every robot’s dreams will improve the performance of all robots.

Big Idea #4: Learning from People

Perception remains one of the most challenging components of robot autonomy. Recently, the ability of large datasets to catalyze perception has proven

to be quite powerful. Large datasets may also have significant utility in planning and control as well.

The online repository of visually recorded objects and human activity is a tremendous resource that robots may soon exploit to improve their ability to understand and interact with the world, including interactions with human beings. Social media sites uploaded more than 1 trillion photos in 2013 and 2014 combined (Meeker 2014, slide 62), and given the growth rate may upload another trillion in 2015. At present, about 300 hours of video are uploaded to the sharing site YouTube every minute, mostly showing people interacting with each other and the environment (see YouTube, <https://www.youtube.com/yt/press/statistics.html>). When 3D sensors become common, the richness of this data store will improve even further.

Most visual information on the Internet is of course not labeled, but clustering techniques can be used to identify similar components in images and videos—for example, in the way that similar faces are grouped now—so that when the meaning of one visual example is learned (perhaps by being labeled by a human), that information can inform the understanding of other images and videos. The utility of this resource is ripe for exponential growth.

Some Implications for the Economy and Workforce

While a Cambrian Explosion in robotics promises to improve the human condition dramatically, it also looms as a disruptive economic force, in part because of its much-discussed potential to make certain human jobs redundant. Yet there is reason to embrace the pending robotics revolution despite such concerns. Consider a Robinson Crusoe economy—a single person stranded on an island with no need for money. Would such a person benefit from a robot, or set of robots, to help with production? Of course. The fundamental economic insight is that robots generate wealth, but the traditional distribution of that wealth through human labor can become problematic. Here, I can offer the thoughts and speculations of a noneconomist on the potential impacts of robots on the ever-evolving labor force. I also discuss some alternatives to the distribution of wealth through labor. These include human services, distribution through capital, and something new—the personal preference information economy.

Output, Demand Satiation, and Human Adaptation

The traditional interaction between technology and the labor market has followed a pattern. Technology enabled an increase in output in certain areas. Demand for the goods produced in that area became at least somewhat satiated, but people were not satiated in their wants and instead soon discovered new areas of demand. Some human labor was displaced as technology expanded, but supply and demand in the labor market drove a series of transitions so that labor shifted to meet the new demand in other areas, and there was no sustained trend to greater unemployment over time. Instead, average wages increased because

technology lifted the productivity of labor. As one example, the Industrial Revolution in textile production created working conditions that were often brutal, but textile output rose and prices fell so that customers (often themselves workers) could afford to buy textiles that would formerly have been out of their economic reach. What consumers didn't spend on textiles could be spent in other sectors, including new sectors. Over time, working conditions in the textile industry improved and labor diversified.

This general pattern has repeated itself in many sectors of the economy over the last two centuries despite warnings every few decades that automation was about to cause mass unemployment. However, this time may be different. When robot capabilities evolve very rapidly, robots may displace a much greater proportion of the workforce in a much shorter time than previous waves of technology. Increased robot capabilities will lower the value of human labor in many sectors. Human abilities as suppliers, even in highly educated societies, evolve slowly. In other words, the increase in robot capabilities may be so rapid that many human workers may find themselves with little to sell.

In the longer run, the diversity and scale of human demand for goods and services has seemed insatiable—so that the labor demanded by the economy did not diminish. But as robot capabilities improve beyond a growing range of human capabilities, will this pattern continue to hold true? One can imagine a future in which many of the material goods that most people want are produced at low cost by advanced robots. Such an economy could evolve in a number of ways. But one possible outcome is that robots may do to many sectors of the economy what the Internet has done to the music business—that is, lead to an economy that pays superstar wages to a small number of exceptionally talented people while paying only a low level of income to many others. In the rest of this section, I explore a few possible ways out of this conundrum.

Human Services

Today, even if a machine-made product is superior in a number of dimensions to a hand-made one, the hand-made product often commands a premium price because it is more difficult to produce and involves the use of a precious commodity—the time of a skilled artisan. As one of many examples, live music continues to be a strong market even as the music recording industry has largely collapsed. This difference exists despite the fact that recorded music is often produced with far higher quality and fidelity than live music, and live music is often played in acoustic environments inferior to those in which music was recorded. However, the excitement of attending a live concert with a crowd of other people is a more valuable experience than listening alone to a recording or the radio. No one proudly wears T-shirts declaring when they listened to a recorded song.

Thus, some human services will probably continue to command a premium compared to robotically produced ones. The question is one of volume: can a future economy be based primarily on personal or hand-crafted goods when close-substitute robotically produced goods become very inexpensive? Concerts are popular, but the vast majority of the music to which people listen is recorded.

The Distribution of Robot Capital Ownership

Imagine a hypothetical economy in which everyone owned a robot and sent their robot to work in their stead. In such a world, the economy could proceed without a hitch, except that we would all have much more leisure time while our robotic stand-ins earned our keep. Of course, the matter of how to initially distribute, trade, and provide safeguards against bankruptcy for robot capital would have to be worked out. But essentially the present system of trading capital, where the intelligence (and significant luck) of investors determines who gets more and who gets less could provide at least some basis for distribution. The late James Albus, an engineer who also served as the head of the Intelligent Systems Division of the Manufacturing Engineering Laboratory at the National Institute of Standards and Technology, explored these issues and possibilities in depth in his 1976 book *Peoples' Capitalism: The Economics of the Robot Revolution*, and its 2011 successor *Path to a Better World: A Plan for Prosperity, Opportunity, and Economic Justice*.

The Personal Preferences Information Economy

In pre-mechanized economies, human beings were born with innate capital for producing economically valuable goods—their bodies. When technology lowered the value of mechanical labor, the economic value of bodies declined but the intrinsic capital value of human brains increased. If brains go the way of bodies, what inherent value will human beings have? Intriguingly, a new inherent human capital has arisen—personal preferences.

Internet companies that had their start producing computer tools like search, email, maps and others have monetized the personal preferences about their users gathered by the tools themselves—which are typically given away “for free.” The gathered information is then sold to advertisers who use it to target individuals most likely to purchase specific goods. The business of these companies is fundamentally the arbitrage of personal preference information. Many people today don't realize the value of their personal preferences, although the substantial profits of the companies that gather and sell such information makes clear its value.

In a future robotic economy, various characteristics of bodies and brains may have much less economic value, but the inherently human value of personal preferences will remain. Were individuals not to surrender personal information so easily, one could imagine Internet companies playing a significant role in wealth distribution by regularly compensating individuals for the value of the information they provide about themselves.

What's Holding Back Robots?

The human brain does much more than store a very large number of associations and access useful memories quickly. It also transforms sensory and other information into generalizable representations invariant to unimportant changes, stores episodic memories, and generalizes learned examples into understanding. The key problems in robot capability yet to be solved are those of generalizable

knowledge representation and of cognition based on that representation. How can computer memories represent knowledge to be retrieved by memory-based methods so that similar but not identical situations will call up the appropriate memories and thoughts?

Significant cues are coming from the expanding understanding of the human brain, with the rate of understanding accelerating because of new brain imaging tools. Some machine learning algorithms, like the Deep Learning approached discussed earlier, are being applied in an attempt to discover generalizable representations automatically. It is not clear how soon this problem will be solved. It may only be a few years until robots take off—or considerably longer. Robots are already making large strides in their abilities, but as the generalizable knowledge representation problem is addressed, the growth of robot capabilities will begin in earnest, and it will likely be explosive. The effects on economic output and human workers are certain to be profound.

References

- Albus, James S.** 1976. *Peoples' Capitalism: The Economics of the Robot Revolution*. New World Books.
- Albus, James S.** 2011. *Path To A Better World: A Plan for Prosperity, Opportunity, and Economic Justice*. iUniverse.
- Bauer, Harald, Jan Veira, and Florian Weig.** 2013. "Moore's Law: Repeal or Renewal?" McKinsey Global Institute, December. http://www.mckinsey.com/insights/high_tech_telecoms_internet/moores_law_repeal_or_renewal.
- Brooks, Rodney.** 2015. In the "Panel Discussion: Artificial Intelligence and the Future of Humans and Robots in the Economy." Malcolm and Carolyn Wiener Annual Lecture, Council on Foreign Relations, February 27, 2015. <http://www.cfr.org/technology-and-science/artificial-intelligence-future-humans-robots-economy/p36197>.
- Cisco.** 2015. "The Zettabyte Era—Trends and Analysis." White Paper, May. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html.
- Hilbert, Martín, and Priscila López.** 2011. "The World's Technological Capacity to Store, Communicate, and Compute Information." *Science*, April 1, 33(6025): 60–65.
- Kuffner, James J.** 2010. "Cloud-Enabled Humanoid Robots." IEEE-RAS International Conference on Humanoid Robotics.
- Li, Fei-Fei.** 2014. "The Digital Sensory System: A Quest for Visual Intelligence in Computers." Presentation at "Stanford Engineering's EngX: The Digital Sensory System" conference, May 20. <https://www.youtube.com/watch?v=yIVsqXzJJqA>.
- Meeker, Mary.** 2014. "Internet Trends 2014—Code Conference." May 28. <http://www.kpcb.com/blog/2014-internet-trends>.
- Miret, Santiago.** 2013. "Storage Wars: Batteries vs. Supercapacitors." Berkeley Energy and Resources Collaborative, November 10. <http://berc.berkeley.edu/storage-wars-batteries-vs-supercapacitors>.
- Moore, Gordon.** 1965. "Cramming More Components onto Integrated Circuits." *Electronics Magazine*, April 19, 38(8).
- Moore, Gordon.** 1975. "Progress in Digital Integrated Electronics." *IEEE, IEDM Tech Digest*, pp. 11–13.
- Parker, Andrew.** 2003. *In the Blink of an Eye: How Vision Sparked the Big Bang of Evolution*. Basic Books.