Intelligence, Wild, Domestic and Artificial

Robert Bercaw

Presented to the Philosophical Club of Cleveland on November 3, 2015

Introduction

Last year Elon Musk, the founder of PayPal, Space-X and Tesla Motors, has called artificial intelligence an existential threat to humanity[i] and likened it to "summoning the demon." He quipped "remember those stories of the guy with the pentagram and the holy water; like you're sure you can control the demon; didn't work out." Stephen Hawkins and Bill Gates have also joined in warning that humans could be superseded by super intelligent machines. These are some pretty impressive intellects, but they left their warnings in the form of generalities; few specifics were provided. I decided to look into this matter for myself, and this is what I found.

The threats to humanity from artificial intelligence and robotics can be placed in three categories: loss of jobs, unintended consequences and AI-robots replacing humans as the dominant species on Earth.

The most immediate threat is that of job loss, which is not unique to AI. It has been a problem throughout the course of the industrial revolution, in which it was largely the story of increasing productivity by replacing human skills and labor with mechanical devices. But, as has often been pointed out, it turned out that the job losses were more than made up by the creation of new jobs, usually at higher skill levels. Today, the computer and the Internet have vastly accelerated the elimination of jobs and have greatly widened the scope of jobs that are being impacted. *Many people are saying it is different this time*; that in the future, the creation of new jobs will not come anywhere close to replacing those that are being lost. Martin Ford, in his recent book,[ii] "The rise of the Robots," makes a strong case that even jobs requiring higher education are subject to replacement.

The replacement of jobs by AI based systems is being accelerated and broadened by the recent development of machine learning, the ability of computers to program themselves. Up till now, computer applications have had to be manually programmed, a manpower intensive and hence expensive process. Jeremy Howard[iii], one of the leading experts in the field, recently outlined the current state of the art; computers are now able to read, understand speech, recognize images in a collection of photos and write sentences with accuracies approaching those of humans. He further asserts that these capabilities are growing at an exponential rate and that computers will soon be on a par with humans—or even surpass them—at these skills, which have essentially defined humanity's uniqueness.

Two well-known examples of machine learning are IBM's Deep Blue and Watson. The first was able to decisively beat Gary Kasperov, the World chess champion. The second trounced the two leading players of the game of Jeopardy, achieving a score greater than the combined scores of the two human contestants.

Unintended consequences

The second danger is that of unintended consequences. As a people, we have a sorry history of seizing on quick fixes to poorly understood problems, only to find that the fixes create even greater problems. Examples range from prohibition, which fostered organized crime, to the introduction of invasive species such as rabbits and cane toads in Australia and kudzu in the U.S.

The increasing delegation of decision making to AI systems, especially those that are self-programmed, raises significant opportunities for new types of unintended consequences. Nick Bostrom's book, "Superintelligence[iv]" argues that AI intelligence could someday surpass human intelligence and then replace humans as the dominant life form on earth. For example, a super-intelligent computer, once given a goal by a human, will spontaneously generate "instrumental" sub goals, which may be impossible to anticipate. To meet the sub goals, the superintelligence might use the web, cellular networks, etc. to acquire unlimited resources or do other mischief.

This is a popular topic in science fiction, in which numerous scenarios have been proposed. It's not clear how real the risks are, but there is a rising concern. Recently an open letter was published[v], signed by Musk, Hawkins and many others, calling for increased research to insure that AI systems are both robust and beneficial to humanity.

A new species

The ultimate unintended consequence would be the creation of a super-intelligent, AI-based, robotic species that would replace our own. Some examples of which have been presented in the movies Her and Ex Machina. Well, just how realistic are scenarios like the ones in these movies and other science fiction. Mostly, the AI species people are talking about correspond roughly to biological ones, and so I'll look at it from that perspective.

The essence of life is that all organisms, from simple bacteria to intelligent beings, continually harvest resources from their environment and use them to maintain themselves, grow and reproduce; anything that doesn't is inanimate. Despite all their impressive capabilities, AI systems are incapable of maintaining or reproducing themselves and hence are inanimate. For example, life support for an AI system is typically provided by a vast infrastructure of power plants, distribution grids and so on. Until an AI creature is able to support itself, it can always be defeated by pulling the plug—like Dave did to HAL in the movie "2001."

Thus, making an AI/robotic system, based on today's computer technology, into a new species requires developing all the technologies required for self-maintenance, growth and reproduction and then moving them into a self-contained unit. I'm not able to assess the details of how hard that would be, but it's clear that it would be an immense task. The alternative would be to develop an AI system based on simple, cell-like modules that are capable of reproduction. But, we don't even know how to make a biological cell, so this alternative is just speculation.

Artificial vs. animal intelligence.

Assuming that an AI species could be created, let's look at how artificial and human intelligence compare. It has often been said that things that are easy for humans are hard for computers and vice versa. The things that are easy for humans are precisely those things that have been supplied to us by evolution to enable us to live and reproduce. And it does appear that those are precisely the things that are difficult for AI systems. Robots are still learning to walk on two legs; their designers haven't a clue about how to make one that could participate in say, a basketball game. It isn't just about making a faster computer; there are fundamental principles yet to be discovered.

AI is good at solving problems where the rules are fairly well defined, even if the execution is complex and ambiguous, such as in Jeopardy, but AI does not appear to be good at looking at issues where formulating the rules is the essence of the problem. No one is expecting AI to develop new scientific theories or great literature.

AI systems are typically used to carry out one or a few applications chosen by their programmers. The brain of a living creature does not have the luxury of choice; it must continually and simultaneously perform all the tasks necessary for the creature's survival.

As animals evolved from invertebrates to fish, to reptiles and finally mammals, their control and information processing needs grew in complexity. Their brains grew apace, but seldom by more than that necessary to meet the needs of the animal's particular lifestyle. The reason is that *neurons are power hogs*, typically consuming something on the order of ten times the calories of other cells in the body. Darwin tells us that a larger brain will only evolve if it provides its animal with a proportionately increased survival rate. For example, predators must have larger brains than herbivores because they need a variety of skills in order to catch their prey. Humans are one of few species to invest in a brain which large with respect to its body size. It has paid off for us. It should be noted, however, that we use 20% of our caloric intake just to maintain our brains, which weigh only 2% of our body mass.

In comparison to the energy efficiency of animal brains, AI systems are way off the chart—the wrong way. For example, Watson consumed several hundred kilowatts of power to defeat its human opponents, whose brains were using just tens of watts--and that's while providing their own life support functions. If we were to imagine a new species using AI for its brain, it would require an improvement in efficiency of four to five orders of magnitude to be competitive with biological species. It's doubtful that this could be achieved by hardware improvements alone.

Job loss

Noting that AI/robotic systems have a very long way to go before a species based on them could be considered a serious threat to humanity, how can we understand why they are so effective in taking jobs away from humans? A small part of it is that their inefficiency isn't a factor, because power is currently very cheap. The larger part, however, is that we have largely structured work in such a way that jobs can easily be automated.

It is generally believed that our species, Homo sapiens, evolved as hunter gatherers under incredibly dangerous and challenging conditions. Our ancestors were slow and lacked the claws and teeth of the major predators; they were definitely in the middle, not the top, of the food chain; their main tool for avoiding predators was their intelligence. Also, securing food required both a deep knowledge of plants, both nutritious and poisonous, and the skills to outwit prey species. Because of the difficulty of a single individual being able to master all these talents, they organized themselves into small social groups or tribes. This, however, required an entirely new set of social skills, and also brought with it a new set of problems, including inter-group strife. Steven Pinker[vi] has compiled a list of archeological studies of 21 prehistoric, hunter-gather and hunter-horticulturists sites. The evidence showed that on average 15% of the deaths in these groups were due to human violence. Today, society is largely organized into states, which have, on average, greatly reduced the rate of death from all forms of violence, including homicides, war, genocide, purges and man-made famines.

What we call civilization began to develop about 10,000 years ago with the invention of agriculture. Agriculture greatly increased food production, but ironically it did not better people's lives. As Malthus explained, population rises exponentially to consume any increases in food production. In fact, the agrarian revolution resulted in both the diets and the variety of activities of most people becoming significantly narrower; the great majority of people ended up living on at a subsistence level. It took the industrial revolution to significantly raise the standard of living, Even now a large proportion of humanity lives on dollar or so a day.

It is important to recognize that our hunter-gatherer ancestors were at least as intelligent, if not more so, than we are today. There is no way to make a direct comparison, but the brain size of humans has actually decreased significantly over the last 10-20 thousand years. Jared Diamond[vii] explains it quite well in his prologue to "Guns, Germs and Steel." He states: "…..In my (33 years of) work with New Guineans, they impressed me as being on the average more intelligent, more alert, more expressive and more interested in things and people around them than the average European or American is." How can this be? We live in a sophisticated, advanced society, while they live in primitive, tribal ones. Diamond sees both genetic and developmental factors involved. Intelligence is a strongly selected attribute in New Guineans, because it is essential to survival in the island's very challenging and dangerous environment, which resembles that of the early hunter-gatherers and hunter-agriculturalists. Meanwhile we have lived for thousands of years in one that is relatively safe, and for most of us, not very challenging. Resistance to infectious diseases, which are common in densely populated regions, is a much larger driver of natural selection than intelligence. As for the development factor, Diamond cites the numerous studies that emphasize the essential role of activity and engagement in promoting the mental development of children. He points out that modern European and American children spend much of their time being passively entertained by television and other media, while New Guinean children spend almost all their time doing things and interacting with others.

Let us define a job as the totality of all the things a person has to do to make a living. If we compare how difficult it would be to create an artificial intelligence to do the job of a hunter-gatherer vs. that of a modern human, it is clear that creating the former would be much more difficult than the latter. Almost everything the hunter-gatherer did was essential to her making a living: i.e., finding food, reproduction

and avoiding death caused by both humans and animals. Many of her tasks had to be done with great urgency and none could be ignored for long.

Civilization has relieved modern humans, and hence their brains, of most of the survival functions that were required for life in the wild. All that is required is that one be able to perform some sort of work that will provide income, and have a modicum of social skills. *Our brains have become domesticated*. The good news is that civilization provides many challenging occupations and has freed our minds to do many discretionary things, if we so choose. The bad news is that our worth has become entwined with the worth of the work that we do.

Many large organizations, wishing to reduce training needs and gain employment flexibility, have defined jobs as narrowly and simply as possible; that is, they have dumbed down the workforce. *The end result is that humanity has only itself to blame for making itself replaceable.*

In summary, it appears as though the threat of AI taking over the world is pretty remote. But their use by organizations is a real threat to the livelihood of many people, so many that it destabilize our current civilization.

This talk was originally prepared for and given to the Transhumanism conference at Juniata College in Pennsylvania. The transhumanism movement that aims to transform the human condition by developing and creating widely available technologies to greatly enhance human intellectual, physical, and psychological capacities. My challenge to them was as follows.

Computers are now being used to make organizations more productive and provide greater profits, but frequently at the expense of the work force. This is the emphasis, because it is the organizations that are paying for the technology. It doesn't have to be that way. Transhumanism would do a great service if it were to provide the means to make individuals more productive than robotic systems and thus remove the incentive to eliminate jobs. People now pay over $100,000 for a college education that will make them more employable, so it stands to reason, that there should be a market for trans-human technology that would provide similar skill augmentations.

Some thoughts along this line can be found in Ray Kurzweil's "How To Create a Mind[viii]" and Clive Thompson's "Smarter Than You Think[ix]."

---

[i] Interview at the MIT AeroAstro Centennial Symposium, Oct 22, 2014
[ii] Martin Ford, "The Rise of the Robots," Basic Books 2015
[iii] Jeremy Howard, "The Wonderful and Terrifying Implications of Computer that Can Learn," TEDxBrussels, Dec 2014
[iv] Nick Bostrom, "SuperIntelligence: Paths, Dangers, Strategies," Oxford 2014
[v] "The Future of Life Open Letter," The Future of Life Institute, Retrieved by Wikipedia, March 4, 2015
[vi] Steven Pinker, "The Better Angels of Our Nature," Viking 2011
[vii] Jared Diamond, "Guns, Germs and Steel," W. W. Norton & Co., 1999, 1997
[viii] Ray Kurzweil, "How To Create a Mind," Viking 2012
[ix] Clive Thompson, "Smarter Than You Think," Penguin 2013