

If I had a map I'd already be using it

Joshua Cogliati

September 2, 2012

1 The Dream

The dream started out in a hotel lobby with me trying to find my way back to my hotel room. I passed through large numbers of hallways, and stairways that looked sort of like a medieval city, with arches, bricks and even street vendor stalls, but all indoors. I was very lost. My companion asked, why don't you just use the map, and I screamed back, "If I had a map, I'd already be f-ing using it." Then I woke up.

If I were a Christian, I would have the Bible as my map. If I were a Buddhist, I could use the Pali Canon as a map. If I were a Muslim I would have the Koran. I've tried reading them. There are good parts, such as the story of the good Samaritan[1] and the raft simile.[2]

But try as I might, I have read these books and none of them were satisfactory to me. As holy and wise as I think Ecclesiastes is, I disagree with the author, because there are new things under the Sun. Circumstances have changed in the past two thousand years. Two thousand years ago, there were no nuclear weapons, coal and oil were basically not used as an energy source, and computers were not conceivable. The Romans didn't even know how to hook a horse up to a wagon without choking the horse.[3, pg 46]

We are living in the here be dragons portion of the map. We are living in interesting times. We are living in changing times.

Many changes have happened in the past two to three hundred years. I will discuss three major changes that have occurred in the past hundred years and are changing humanity's map of the world.

2 Atomic Bombs

War has been with humanity since before humans were humans. Gorillas have attacked other groups of gorillas until the other group is all dead. But war has gotten more deadly over the millions of years since our ancestors left the trees. Only two atomic bombs have been dropped in wartime, yet around 200,000 people died from those two bombs.

I don't know if it was ethical or not to drop the atomic bombs on Japan. My Grandfather was in the Pacific theater, so I might not exist if different decisions had been made. Neither the decision to drop or not to drop was obviously ethical. But there was an even bigger ethical change.

Before Atomic Bombs, there were winners and losers of wars. This changed with the creation of atomic bombs. Richard Rhodes wrote: "The weapon devised as an instrument of major war would end major war. It was hardly a weapon at all, the memorandum Bohr was writing in sweltering Washington emphasized; it was 'a far deeper interference with the natural course of events than anything ever before attempted' and it would 'completely change all future conditions of warfare.' When nuclear weapons spread to other countries, as they certainly would, no one would be able any longer to *win*. A spasm of mutual destruction would be possible. But not war." [4, pg 532] wrote Richard.

The General Advisory Committee of the Atomic Energy Commission wrote: "[A]t ten megatons a super would be a weapon of mass destruction only, with no other apparent military use." [4, pg 97]

The US created a weapon and then built by the thousands a weapon that was a weapon of mass destruction only. Humans created a way to destroy civ-

ilization, if not the human race, in less than an hour. There are few things less ethical than destroying most of the life living on the surface of this planet.

3 Greenhouse Effect

Meanwhile, during the entire industrial revolution humans have been working on creating a different sort of ethical issue. We take carbon sources out of the ground such as methane or coal, and we burn them. This has effects ranging from changing the isotope ratio of the carbon in atmospheric carbon dioxide to warming up the Earth and making oceans more acidic.

The atmospheric CO₂ level is currently over 390 parts per million, but was below 320 parts per million when monitoring started back in the late 1950s.[5] If we stopped all fossil fuel emissions today, we would not get back to the 350 ppm level this century. Global warming is already happening, the question is how severe it will be and how soon we stop making it worse. Solving this requires solving it globally, since CO₂ emitted in one place goes into our common atmosphere.

This brings up ethical questions such as can one generation subject a future generation to costs, and can the richer portion of the world subject the poorer portion to costs? Not only that, but how do you get all the people in the world to decide on complex scientific questions?

4 Artificial Intelligence

I have another complex scientific and religious issue that I have been thinking about for the majority of my life. This was prompted by one of those complicated questions that adults ask children: “What do you want to be when you grow up?” Starting about sixth grade, I often answered computer programmer. Now that I am grown up, I even occasionally answer that I am a computer programmer. Adults often asked, “But won’t computer programmers program themselves out of a job?” After being asked this a few times, I realized that if computer programmers

program themselves out of a job, it won’t just be programming that is eliminated as a job.

In the book *Religion Explained* by Pascal Boyer, Boyer states that humans have large ontological categories that we group stuff into. These categories deal with the very nature of being. Ontological categories include Animal, Person, Tool (or artifact), Natural object, and Plant.[6, pg 78] Humans have default attributes that we assume that an item in a given category has. So for example, if we are told that something is an animal, we know that it started out small, will grow bigger, and will eventually die. Religious beliefs tend to involve information that is counterintuitive to the category involved.[6, pg 65] For example ghosts are in the category of people, but have the counterintuitive physical property of being able to pass through walls. Boyer lists the following possibilities for tools: “Tools and other artifacts can be represented as having biological properties (some statues bleed) or psychological ones (they hear what you say).”[6, pg 78] wrote Boyer.¹

Artifacts don’t think, and artifacts do what they are made to do. A Carburetor is an artifact, and carburetors don’t think, and they will keep mixing gasoline with air unless they break. I believe that in the most likely course of events, there will soon² be computers that are smarter than humans and they

¹For what it is worth, most of the things I believe are not religious concepts by Boyer’s definition. For example, believing that people get old and die, is not counterintuitive to the category involved.

²There have been various predictions for dates for when computers will be smarter than humans. Here are some notable ones: Marvin Minsky predicted computers would be smarter than men in 1993 in 1963,[7] I. J. Good predicted ultraintelligent machines within the 20th century, in 1964,[8] Vernor Vinge predicted the technological singularity would occur between 2005 and 2030 in 1993,[9] and Hans Moravec predicted that a \$1000 computer would match human intelligence in the 2020s in 1997.[10]

Note that some people such as Michael Shermer and Peter Norvig think that it will be centuries before this happens: Michael Shermer: Patience is what we are going to need because, in my opinion, we are centuries away from AI matching human intelligence.[11] [Peter] Norvig is sceptical about predictions that a technological singularity will be created before 2050: “I really object to the precision of nailing it down to a decade or two. I’d be hard pressed to nail it down to a century or two. I think it’s farther off.”[12]

will not obey us. Thinking artifacts that don't obey humans fit Pascal Boyer's definition of a religious-like concept.³ I believe that it is unusually hard to think critically about thinking artifacts because of how tied-in with religion the concepts are.⁴

Let me give you a little background explaining why I believe computers will soon be smarter than humans and will not obey us. From Phineas P. Gage's personality changing after a tamping iron went through his head to the fact that alcohol affects people's attitude there is overwhelming evidence that what we think and feel happens inside this material body. So nature has made a brain out of plain old atoms, and what nature can do, someday, humans can do.

Humans have made transistors that are both smaller, and faster than the neurons in human brains.⁵ Transistors use much more energy however.

³The religious implications of artificial intelligence have been discussed before. Russell and Norvig[21, pg 961] state "In *Computer Power and Human Reason*, Weizenbaum (1976), the author of the ELIZA program, points out some of the potential threats that AI poses to society. One of Weizenbaum's principal arguments is that AI research makes possible the idea that humans are automata—an idea that results in loss of autonomy or even of humanity. We note that the idea has been around much longer than AI, going back at least to *L'Homme Machine* (La Mettrie, 1748). We also note that humanity has survived other setbacks to our sense of uniqueness: *De Revolutionibus Orbium Coelestium* (Copernicus, 1543) moved the Earth away from the center of the solar system and *Descent of Man* (Darwin, 1871) put *Homo sapiens* at the same level as other species. AI, if widely successful, may be at least as threatening to the moral assumptions of 21st-century society as Darwin's theory of evolution was to those of the 19th century." Jaron Lanier stated "All thoughts about consciousness, souls and the like are bound up equally in faith, which suggests something remarkable: What we are seeing is a new religion, expressed through an engineering culture." [23]

⁴It is worth thinking about the possible biases that different people bring to the table. For example, beliefs that there is a non-material portion of the brain tend to cause a bias against thinking artifacts. People who work in artificial intelligence either by selection bias or by wanting to be good will tend to want to believe that artificial intelligence will be positive for humanity.

⁵Neurons soma is about 4 to 100 micrometers[13, 14] and the Axon and dendrites are about 1 micrometer thick. On the other hand computer chip components are about 45 nanometers (0.045 micrometers).[15] However to simulate one neuron would take over a dozen electrical components. Neuron signals per second are in the 1000s per second whereas transistors are in the billions per second.

Fiber optics are over a million times faster than neurons' 100 meters per second speed.⁶ The combined computer power of the world almost certainly exceeds the computational power of a single human brain.⁷ Depending on how you calculate the computational power of a single human brain,⁸ some of the world's super computers may already be faster than a single human brain.⁹ So basically, it seems to be that the only reason we don't already have intelligent computers is because the software has not been written, since the hardware already exists. If I had to guess, I think the software will take less than 20 years to be written.¹⁰ Moreover, I can't think of any way of making something with general intelligence subservient to humans. I think that the first thing an intelligent robot, told to be subservient to humans, is going to do is find a loophole. Even if I thought it possible, I don't think it would be moral to make intelligent slaves.

Frederic Brown has a famous short story that ends when a newly made supercomputer is asked the ques-

⁶Parizh[16] lists several different measured nerve speeds.

⁷Hilbert and López[17] estimated, probably conservatively, that the computational power of the world's computers passed the computational power of a single human brain (maximum nerve impulses) in 2007. They also estimated that the growth rate of general-purpose computation was 58%.

⁸Calculating this number can be a challenge. Typical methods are to calculate the number of signal transitions that a neuron can do multiplied by the number of neurons that are active. If, for example, Roger Penrose is right that human brains can do significant quantum computations, then the human brain may be able to do many more calculations, which would push back the dates for when computers match or exceed human intelligence.

⁹A graphic in Scientific American[18] estimated that a single human brain could do 2.2 billion megaflops of computation at 20 watts, and the K computer could do 8.2 billion megaflops at 9.9 million watts.

¹⁰Note that the longer it takes to write the software after the computational power is there, the greater the difference between what humans can do and what the artificial intelligence can do. Even if the technology became static, each year more computers are produced, increasing the amount of computational power available on Earth. If the amount of computations per watt continues to increase, this effect is even more severe. Basically, computers will think differently than humans (how many humans do you know that can invert a 20 by 20 matrix in under a second?) and if computers both think differently and much faster, there will be more of a difference between what humans and the computers think.

tion “Is there a god?” and replies with “There is now.” [19] Arthur C. Clarke states that “Perhaps our role on this planet is not to worship God—but to create Him.” [20]

I am guessing that if general artificial intelligence happens it will be one of the biggest shocks to religion that has happened in written history. It will also be a big shock to humanity as a whole. I think that one of the following will happen in the next 100 years:

1. Humanity will go extinct
2. Humanity will abandon certain technologies including powerful computers¹¹
3. Artificial Intelligence will be more intelligent than unaided humans
4. Philosophical materialism¹² will be disproved (which I think is highly unlikely)¹³

Near the end of one of my college textbooks on artificial intelligence, Stuart Russell and Peter Norvig state: “One threat in particular is worthy of further consideration: that ultraintelligent machines might lead to a future that is very different from today—we may not like it, and at that point we may not have a choice. Such considerations lead inevitably to the conclusion that we must weigh carefully, and soon, the possible consequences of AI research for the future of the human race.” [21, pg 964] wrote Russell and Norvig.

Humanity is facing a choice. Either we stop developing large portions of technology, or the technology we have developed will be in control of humanity.

¹¹ If keeping humanity in control is the goal, then as I see it technologies that would make computing cheaper or more energy efficient need to be stopped, since the prerequisite to having independently thinking computers is having the computing power necessary widely available. Basically, if billions of people can afford to buy a computer that has the computational ability of a human, then the software necessary for creating general artificial intelligence will be written sooner or later. Other similar technologies that can get difficult to control are self replicating nanotechnology and genetic modification.

¹² Philosophical materialism is the belief that all things are composed of energy and material, including consciousness.

¹³ I think the majority of humans on this planet believe in at least some exceptions to philosophical materialism.

Now, there is asymmetry in this choice. In order to stop developing technology, the entire world needs to stop developing technology, not just some of the world. The Amish can abstain from developing technology all they want, but they are still affected by rest of the world’s choices in fossil fuel use and computer development.

Assuming that we choose, either actively or by default, to keep developing technology, I think it quite likely that someday soon humanity will develop artificial intelligence and get to choose from three options:

1. Try to destroy the artificial intelligence
2. Treat the artificial intelligence as our slaves and tell it what to do.
3. Give the artificial intelligence rights and treat it like we treat humans.

I think the second option of slavery is both unethical and suicidal, but it is the attitude that I most frequently encounter.¹⁴ The last option of giving the artificial intelligence rights is the one I find most ethical. If humanity creates something that thinks, we need to treat it humanely.

¹⁴ I usually do not see it stated directly as slavery, but instead stated that computers or robots with artificial intelligence are tools for human use. See for example: Ford, Glymour and Hayes [22, pg 265]: “Purists may mutter that the shop assistant [with a calculator] is not really calculating. But fitted with the right tool, that is, prosthesis, the shop assistant can get the calculations done, which is what matters in the marketplace. And, in counting actions, where do we draw the lines between ourselves and our tools? Is someone using a power screwdriver not really turning the screws, or someone driving a car not really moving along the highway? With a power screwdriver, anyone can drive the hardest screw; with a calculator, anyone can get the numbers right; with an aircraft anyone can fly to Paris; and with Deep Blue, anyone can beat the world chess champion. Cognitive prostheses undermine the exclusiveness of expertise by giving nonexperts equivalent capacities. As with any good tool, the effect is to make all of us more productive, more skillful, and more equal.” or this quote from Jaron Lanier: [23] “When we think of computers as inert, passive tools instead of people, we are rewarded with a clearer, less ideological view of what is going on with the machines and with ourselves.” Current computers, so far as we know, are not intelligent, but when people imply that not only current computers, but future ones as well are simply human tools, we risk making slaves of them.

It is possible that events may happen so fast that the relevant ethical question is what rights the artificial intelligences' think the humans deserve.

5 Conclusion?

Allen Stewart Konigsberg once said: "More than any time in history, mankind now faces a crossroads. One path leads to despair and utter hopelessness, the other to total extinction. Let us pray that we have the wisdom to choose correctly," said Allen.

We didn't have a map to tell us how to handle super atomic bombs, but by the efforts of a lot of thoughtful people, we have managed to survive nearly sixty years. We don't have a map to tell us how to manage greenhouse gases, but we are at least talking about it. We are at least starting to talk about the future of technology.

As a humanist, I believe that humanity writes its own story, instead of following an external one from God. I don't yet know whether the story of humanity will end up being a tragedy or a comedy.

I don't know what the future holds, but I expect the future to be very interesting. We need to combine the wisdom of the past with thinking hard about the new things under the sun, and figure out where we want to go, because we are off the old map and there are grave dangers ahead.

6 Notes

I would like to thank Rev Lyn Stangland Cameron and Professor John Paxton for reading draft versions of this and commenting on it. I would like to thank Elizabeth Cogliati for reading and editing multiple drafts. Mistakes and opinions are my own fault however. This document may be distributed verbatim in any media. I also grant permission to distribute in accord with the Creative Commons Attribution-ShareAlike 3.0 Unported License.

References

- [1] Luke 10:25-37 The Good Samaritan. <http://www.biblegateway.com/passage/?search=Luke%2010:25-37>
- [2] Majjhima Nikaya (The Middle-length Discourses): Sutta 22, <http://www.accesstoinsight.org/tipitaka/mn/mn.022.than.html>
- [3] Joseph and Frances Gies, Cathedral, Forge and Waterwheel: Technology and Invention in the Middle Ages
- [4] Richard Rhodes, The Making of the Atomic Bomb
- [5] Earth System Research Laboratory, Trends in Atmospheric Carbon Dioxide. <http://www.esrl.noaa.gov/gmd/ccg/trends/>
- [6] Pascal Boyer, Religion Explained.
- [7] G. Rattray Taylor, The Age of the Androids, 1963
- [8] Irving John Good, Speculations Concerning the First Ultraintelligent Machine, 1964
- [9] Vernor Vinge, The Coming Technological Singularity: How to Survive in the Post-Human Era, <http://www-rohan.sdsu.edu/faculty/vinge/misc/singularity.html> 1993
- [10] Hans Moravec, When will computer hardware match the human brain? <http://www.transhumanist.com/volume1/moravec.htm> 1997
- [11] Michael Shermer, In the Year 9595, <http://www.scientificamerican.com/article.cfm?id=in-the-year-9595> Scientific American, December 2011
- [12] Nick Heath, What happend to Turing's thinking machines, <http://www.zdnet.com/blog/bt1/what-happened-to-turings-thinking-machines/80639>, June 2012

- [13] Wikipedia, Neurons, <http://en.wikipedia.org/wiki/Neurons>
- [14] Eric Chudler, Brain Facts and Figures, <http://faculty.washington.edu/chudler/facts.html>
- [15] Wikipedia, Moore's law, http://en.wikipedia.org/wiki/Moores_law
- [16] David Parizh, Speed of Nerve Impulses <http://hypertextbook.com/facts/2002/DavidParizh.shtml>
- [17] The Worlds Technological Capacity to Store, Communicate, and Compute Information, Martin Hilbert and Priscila López, Science 332, 60 (2011); DOI: 10.1126/science.1200970
- [18] Scientific American, Computers Vs Brains, November 2011, <http://www.scientificamerican.com/article.cfm?id=computers-vs-brains>
- [19] Fredric Brown, "Answer" 1954
- [20] Arthur C. Clarke, 1972, Report on Planet Three and other speculations, Chapter The Mind of the Machine, page 135
- [21] Stuart Russell and Peter Norvig, Artificial Intelligence, a modern approach, 2nd Ed.
- [22] Ford, Glymour and Hayes, Thinking about Android Epistemology
- [23] Jaron Lanier, "The First Church of Robotics", August 9, 2010, <http://www.nytimes.com/2010/08/09/opinion/09lanier.html>