

Intelligence is not Artificial

*Why the Singularity is not Coming any Time Soon And
Other Meditations on the Post-Human Condition and
the Future of Intelligence*

Piero Scaruffi

ISBN 978-0-9765531-9-9

All Rights Reserved © 2015 by Piero Scaruffi

No part of this book may be reproduced or transmitted in any form or by any means, graphic, electronic or mechanical, including photocopying, recording, taping or by any information storage retrieval system, without the written permission of the author (<http://www.scaruffi.com>). Your support of the authors' rights is appreciated.

Table of Contents

- What this Book is about
- Sociological Background
- A Brief History of Artificial Intelligence
- Footnotes in the History of Artificial Intelligence
- A Premise to the History of Artificial Intelligence
- Brain Simulation and Intelligence
- The Body
- Intermezzo: We may Overestimate Brains
- Childhood
- The State of the Art in Artificial Intelligence
- An Easy Science
- Intermezzo and Trivia: the Original App
- Don't be Fooled by the Robot
- A Consumer's Rant Against the Stupidity of Machines: Reverting Evolution?
- The Singularity as the Outcome of Exponential Progress
- A Look at the Evidence: A Comparative History of Accelerating Progress
- Intermezzo: In Defense of Regress
- The Feasibility of Superhuman Intelligence
- What is the Opposite of the Singularity?
- Intermezzo: The Attention Span
- Semantics
- Non-human Intelligence is Already Here
- From Non-human to Super-human Intelligence
- The Accelerated Evolution of Machines
- Human Obsolescence
- Historical Intermezzo: Jobs in the 2000s
- Marketing and Fashion
- The Audience of the Singularity
- In Defense of Progress

- Intermezzo: You Are a Gadget
- Task-specific vs General-purpose Intelligence
- Intelligent Behavior from Structured Environments
- Intermezzo: Will Intelligent Machines Return to Chaotic Environments?
- The Multiplication of Appliances and Artificial Intelligence by Enumeration
- Mind Uploading and Digital Immortality
- Machine Immortality and the Cloud
- Corollary: Digital Media Immortality
- Ethical Intermezzo: The Moral Consequences of Human Immortality
- Philosophical Intermezzo: Do We Really Want Intelligence at All?
- Analog vs Digital
- Teaser: Machine Ethics
- Summarizing and The Timeframe for Artificial General Intelligence
- Intermezzo: We may Overestimate Human Brains
- The Real Future of Computing
- The Future of Miniaturization
- The Moral Issue: Who's Responsible for a Machine's Action?
- The Dangers of Machine Intelligence: Machine Credibility
- The Dangers of Machine Intelligence: Speed Limits for Machines?
- The Dangers of Machine Intelligence: Criminalizing Common Sense
- The Dangers of Machine Intelligence: You Are a Budget
- Why We Shouldn't Talk About the Singularity
- The Future of Human Intelligence
- A Conclusion: Sociology Again
- P.S. What i did not Say
- Appendix: The Dangers of Clouding - Wikipedia as a Force for Evil
- Appendix: The Myth of Longevity
- Appendix: The Medium is the Brain

- Appendix: The Era of Objects
- Appendix: I Ran Out of Space
- A Timeline of Neuroscience
- A Timeline of Artificial Intelligence
- Readings on the Singularity
- Alphabetical Index

What this Book is about

Writers, inventors and entrepreneurs, impressed by progress in several scientific fields, are debating whether we may be heading for a "singularity" in which machines with super-human intelligence will arise and multiply. In parallel enthusiastic coverage in the media has widely publicized machines performing sophisticated tasks, from beating the world's chess champion to driving a car, from recognizing animals in videos to beating human experts on quiz shows. These stories have reignited interest in the discipline of Artificial Intelligence, whose goal is to create machines that are as intelligent as humans.

First of all, this book provides a "reality check" of sorts on simulating human intelligence and achieving superhuman intelligence. I show that, in a society driven by media that desperately need sensational news to make money and in an academic world increasingly driven by the desire to turn research into Silicon Valley start-ups, technological progress in general, and progress in computer science in particular, is often overrated. I wanted to dispel some notions, and my version of the facts may sound controversial until you read my explanations. Progress in Artificial Intelligence since its founding has been negligible, and one reason is, ironically, that computers have become so much more (computationally) powerful.

I also wanted to caution that we tend to exaggerate the uniqueness of our age, just as previous generations have done. The very premise of the singularity theory is that progress is accelerating like never before. I argue that there have been other eras of accelerating progress, and it is debatable if ours is truly so special. The less you know about the past the more likely you are to be amazed by the present.

There is certainly a lot of change in our era. But change is not necessarily progress, or, at least, it is not necessarily progress for everybody. Disruptive innovation is frequently more about disruption

than about innovation because disruption creates huge new markets for the consumer electronics industry.

There is also an excessive metaphysical belief that human intelligence is the current apex of evolution. Maybe so, but it is worth cautioning that non-human intelligence is already among us, and is multiplying rapidly, but it is not a machine: countless animals are capable of feats that elude the smartest humans. For a long time we have also had machines capable of performing "superhuman" tasks. Think of the clock, invented almost 1,000 years ago, that can do something that no human can do: it can tell how many hours, minutes and seconds elapse between two events.

Once we realize that non-human intelligence has always been around, and that we already built super-human machines centuries ago, the discussion can be framed in more historically and biologically meaningful terms.

The new generations missed the debates of the previous decades (the "Turing test", the "ghost in the machine", the "Chinese room", etc); and some of us think that these new generations, trained in hyper-specialized disciplines, don't have the knowledge to understand them even if they were forced to read about them. Therefore it is much easier for the new A.I. practitioners to impress the new generations. I have summarized the various philosophical arguments in favor of and against the feasibility of machine intelligence in my book "The Nature of Consciousness" and i won't repeat them here. I will, however, at least caution the new generations that i "grew up" (as far as cognitive science goes) at a time when the term "intelligence" was not "cool" at all: too vague, too unscientific, too abused in popular literature to lend itself to scientific investigation.

In fact, the mother of all problems in this debate is at the very source: "singularity" and "superhuman intelligence" are non-scientific terms based on non-scientific chatting. The term "intelligence" itself is hardly scientific. Ask one hundred psychologists and you will get one hundred different definitions. Ask

philosophers and you will get thick tomes written in a cryptic language. Ask neurobiologists and they may simply ignore you.

I also feel that this discussion should be complemented with an important (more important?) discussion about the changes in human intelligence due to the increased "intelligence" of machines. This change in human intelligence may have a stronger impact on the future of human civilization than the improvements in machine intelligence. To wit: the program of turning machines into humans is not very successful yet, but the program of turning humans into machines (via an endless repertory of rules and regulations) is very successful.

Finally, i am intrigued by another sociological/anthropological aspect of this discussion: humans seem to have a genetic propensity to believe in higher forms of intelligence (gods, saints, UFOs, ...) and the myth of the Singularity could simply be its manifestation in our post-religious 21st century.

Piero Scaruffi, September 2013 & June 2015

P.S. Since the Singularity narrative is often identified with Ray Kurzweil, i want to dispel the notion that this book could be read as a direct attack against him. While i disagree with some of his optimistic views, i feel great admiration for Ray Kurzweil, one of the few scientists who actually makes predictions that can be tested.

P.S.: Yes, i don't like to capitalize the first person pronoun "i".

Sociological Background

Historians, scientists and poets alike have written that the human being strives for the infinite. In the old days this meant (more or less) that s/he strives to become one with the god who created and rules the world. As atheism began to make strides, Arthur Schopenhauer rephrased the concept as a "will to power". Friedrich Nietzsche confirmed that god is dead, and the search for "infinite" became a mathematical and scientific program instead of a mystical one. Bertrand Russell, David Hilbert and others started a logical program that basically aimed at making it easy to prove and discover everything that can be. The perspective therefore changed: instead of something that humans have to attain, the infinite has become something that humans will build.

One of the consequences of this line of research was the creation of the digital electronic computer, the physical implementation of a thought experiment by Alan Turing. The same mathematician wrote a pioneering paper about machine intelligence and a few years later the term "artificial intelligence" was already popular among both scientists and philosophers.

The idea behind the Singularity, a concept popularized by Ray Kurzweil's "The Age of Intelligent Machines" (1990) and by his subsequent, highly-successful public-relationship campaign, is that we are about to witness the advent of machines that are more intelligent than humans, so intelligent that humans can neither control them nor understand them.

The discipline of Artificial Intelligence, that had largely languished, has staged a revival of sorts, at least in the eyes of public opinion. Achievements in the field of A.I. are often hailed by the mainstream media as steps towards machine domination. (In this book i capitalize Artificial Intelligence when i am referring to the discipline, while using lowercase "artificial intelligence" to refer to an intelligent machine or an intelligent software).

In the age that has seen the end of human space exploration, the retirement of the only commercial supersonic airplane, the decline of nuclear power, and the commercialization of the Internet (an event that basically turned a powerful scientific tool into little more than a marketing tool and a form of light entertainment), machine intelligence seems to bring some kind of collective reassurance that we are not, after all, entering a new Dark Age; on the contrary, we are witnessing the dawn of a superhuman era. Of course, decades of science fiction books and movies helped create the ideal audience for this kind of scenario.

However, the tone and the (very weak) arguments in favor of the Singularity do remind one of religious prophecies, except that this time the coming prophet will arise from our own doing instead of being sent by an external divinity. In a sense, this is a religion according to which we are creating the divinity.

The idea of the Singularity is fascinating because it plays the history or religion backwards. Religion traditionally is meant to explain the mystery of the complexity of the universe, the miracle of life, the purpose of consciousness. Even some of today's eminent scientists subscribe to the "creationist" notion that a superhuman intelligence was required to create the world. This theory is frequently called "intelligent design" but it would be more appropriate to call it "super-intelligent design" if "intelligent" refers to human intelligence. The whole point of religion is precisely to explain something that exists and that human intelligence could never possibly build. The hidden assumption of religion is that all the laws of nature that humans can possibly discover will never be enough to explain the mysteries of the universe, of life, of the soul. Whatever can be explained by those mathematical laws can also be implemented by humans, and therefore does not require the existence of supernatural forces. God is a singularity, the singularity that preceded human intelligence and that created it.

The hypothesis of a coming singularity due to super-intelligent machines provides a mirror image of this story. The original singularity (God) was needed to explain the unexplainable. The new

singularity (the machine) will be unexplainable. Human intelligence cannot understand the workings of the God of the past that created human intelligence; nor can it understand the workings of the super-intelligent machine of the future that human intelligence will have created.

It is sometimes difficult to argue with the Singularity crowd because they often seem unaware that some of the topics they discuss have been discussed for a long time, with pros and cons, by philosophers and scientists. In its worst manifestation the Singularity movement is becoming the religion of high-tech nerds who did not study history, philosophy, science and not even computer science. At its best, however, it helps acquaint the general public with a society of (software and hardware) robots that is inevitably coming, although its imminence might be wildly overstated.

It may not be a coincidence that the boom of interest in the Singularity originated in the USA (a country well acquainted with apocalyptic preachers, conspiracy theories and UFO sightings) and that it originated after the year 2,000. The new millennium triggered a revival of interest in the likes of Nostradamus, and gave us Harold Camping's Biblical calculations that the end of the world was coming on October 21 of 2011 and the theory that the end of the Mayan calendar (December 21 2012) marked the end of the world.

A Brief History of Artificial Intelligence

One can start way back in the past with the ancient Greek and Chinese automata of two thousand years ago or with the first electromechanical machines of a century ago, but to me a history of machine intelligence begins in earnest with Alan Turing's "universal machine", originally conceived in 1936. He did not personally build it, but Turing realized that one could create the perfect mathematician by simulating the way logical problems are solved: by manipulating symbols. The first computers were not Universal Turing Machines (UTM), but every computer built since the ENIAC

(1946), including all the laptops and smartphones that are available today, is a UTM. Because it was founded on predicate logic, which only admits two values ("true" and "false"), the computer at the heart of any "intelligent" machine relies on binary logic (ones and zeroes).

Cybernetics (that can be dated back to the 1943 paper "Behavior, Purpose and Teleology" co-written by mathematician Norbert Wiener, physiologist Arturo Rosenblueth and engineer Julian Bigelow) did much to show the relationship between machines and (non-intelligent) living organisms. One can argue that machines are form of life or, viceversa, that living organisms are forms of machinery.

However, "intelligence" is commonly considered one or many steps above the merely "alive": humans are generally considered intelligent (by fellow humans), whereas worms are not.

The "Turing Test", introduced by the same Alan Turing in his paper "Computing Machinery and Intelligence" (1950), has often been presented as the kind of validation that a machine has to pass in order to be considered "intelligent": if a human observer, asking all sorts of questions, cannot tell whether the agent providing the answers is human or mechanical, then the machine has become intelligent (or, better, as intelligent as the human being). The practitioners of Artificial Intelligence quickly split in two fields. One, pioneered by Allen Newell and Herbert Simon with their "Logic Theorist" (1956), basically understood intelligence as the pinnacle of mathematical logic, and focused on symbolic processing. The first breakthrough in this branch of A.I. was probably John McCarthy's article "Programs with Common Sense" (1959): McCarthy understood that some day machines would easily be better than humans at many repetitive and computational tasks, but "common sense" is what really makes someone "intelligent" and common sense comes from knowledge of the world. That article spawned the discipline of "knowledge representation": how can a machine learn about the world and use that knowledge to make inferences. This approach was somehow "justified" by the idea

introduced by Noam Chomsky in "Syntactic Structures" (1957) that language competence is due to some grammatical rules that express which sentences are correct in a language. The grammatical rules express "knowledge" of how a language works, and, once you have that knowledge (and a vocabulary), you can produce any sentence in that language, including sentences you have never heard or read before.

The rapid development of computer programming helped this field take off, as computers were getting better and better at processing symbols: knowledge was represented in symbolic structures and "reasoning" was reduced to a matter of processing symbolic expressions. This line of research led to "knowledge-based systems" (or "expert systems"), such as Ed Feigenbaum's Dendral (1965), that consisted of an inference engine (the repertory of legitimate reasoning techniques) and a knowledge base (the "common sense" knowledge). This technology relied on acquiring knowledge from domain experts in order to create "clones" of such experts (machines that performed as well as the human experts). The limitation of expert systems was that they were "intelligent" only in one specific domain.

Meanwhile, the other branch of Artificial Intelligence was pursuing a rather different approach: simulate what the brain does at the physical level of neurons. Since neuroscience was just in its infancy (medical machines to study living brains would not become available until the 1970s), computer scientists only knew that the brain consists of a huge number of interconnected neurons, and neuroscientists were becoming ever more convinced that "intelligence" was due to the connections, not to the individual neurons. A brain can be viewed as a network of interconnected nodes, and our mental life as due to the way messages travel through those connections from the neurons of the sensory system up to the neurons that process those sensory data and eventually down to the neurons that generate action. The neural connections can vary in strength from zero to infinite. Change the strength of the neural connections and you change the outcome. In other words,

the strength of the connections can be tweaked to cause different outputs for the same inputs: the problem for those designing "neural networks" consists in fine-tuning the connections so that the network as a whole comes up with the correct interpretation of the input; e.g. with the word "apple" when the image of an apple is presented. This is called "training the network". For example, showing many apples to the system and forcing the answer "APPLE" should result in the network adjusting those connections to recognize apples. Hence the alternative term "connectionism" for this branch of A.I. Frank Rosenblatt's Perceptron (1957) and Oliver Selfridge's Pandemonium (1958) were the pioneer "neural networks": not knowledge representation and logical inference, but pattern propagation and automatic learning. Compared with expert systems, neural networks are dynamic systems (their configuration changes as they are used) and predisposed to learning by themselves (they can adjust their configuration). "Unsupervised" networks, in particular, can discover categories by themselves (e.g., discover that several images refer to the same kind of object, like an apple).

There are two ways to solve a crime. One way is to hire the smartest detective in the world, who will use experience and logic to find out who did it. On the other hand, if we had enough surveillance cameras placed around the area, we would scan their tapes and look for suspicious actions. Both ways may lead to the same conclusion, but one uses a logic-driven approach (symbolic processing) and the other one uses a data-driven approach (ultimately, the visual system, which is a connectionist system).

In 1969 Marvin Minsky and Samuel Papert published a devastating critique of neural networks (titled "Perceptrons") that virtually killed the discipline. At the same time expert systems were beginning to make inroads at least in academia, notably Bruce Buchanan's Mycin (1972) for medical diagnosis and John McDermott's Xcon (1980) for product configuration, and, by the 1980s, also in the industrial and financial worlds at large, thanks especially to many innovations in knowledge representation (Ross

Quillian's semantic networks, Minsky's frames, Roger Schank's scripts, Barbara Hayes-Roth's blackboards, etc). Intellicorp, the first major start-up for Artificial Intelligence, was founded in Silicon Valley in 1980.

One factor that certainly helped the symbolic-processing approach and condemned the connectionist approach was that the latter uses complex algorithms, i.e. it requires computational power that at the time was rare and expensive.

(Personal biography: i entered the field in 1985 and went on to lead the Silicon Valley-based Artificial Intelligence Center of the largest European computer manufacturer, Olivetti, and i later worked for Intellicorp).

Footnotes in the History of Artificial Intelligence

There were many side tracks that didn't become as popular as expert systems and neural networks. Stanford Research Institute's robot Shakey (1969) was the vanguard of autonomous vehicles. IBM's Shoebox (1964) debuted speech recognition. Conversational agents such as Joe Weizenbaum's Eliza (1966) and Terry Winograd's Shrdlu (1972) were the first practical implementations of natural language processing. In 1968 Peter Toma founded Systran to commercialize machine-translation systems. Refining an idea pioneered by Ingo Rechenberg in his thesis "Evolution Strategies" (1971), John Holland introduced a different way to construct programs by using "genetic algorithms" (1975), the software equivalent of the rules used by biological evolution: instead of writing a program to solve a problem, let a population of programs evolve (according to those genetic algorithms) to become more and more "fit" (capable of finding solutions to that problem). In 1976 Richard Laing introduced the paradigm of self-replication by self-inspection ("Automaton Models of Reproduction by Self-inspection", 1976) that 27 years later would be employed by Jackrit Suthakorn and Gregory Chirikjian at Johns Hopkins University to build a rudimentary self-replicating robot ("An autonomous self-replicating

robotic system", 2003). Since 1979 Cordell Green experimented with automatic programming, software that can write software the same way a software engineer does. In 1990 Carver Mead described a "neuromorphic" processor, a processor that emulates the human brain.

A Premise to the History of Artificial Intelligence

Surprisingly few people ask "why?" Why did the whole program of A.I. get started in the first place? What is the goal? Why try and build a machine that behaves (and feels?) like a human being?

There were and there are several motivations. I believe the very first spark was pure scientific curiosity. A century ago an influential German mathematician, David Hilbert, outlined a program to axiomatize mathematics as a sort of challenge for the world's mathematicians. In a sense, he asked if we can discover a procedure that will allow anybody to solve any mathematical problem: run that procedure and it will prove any theorem. In 1931 Kurt Goedel proved his theorem of incompleteness, which basically answered "No, that's not possible, because there will always be at least one proposition that we cannot prove true or false"; but in 1936 Alan Turing offered his solution, now known as the Universal Turing Machine, which is as close as we can get to Hilbert's dream procedure. Today's computers, including your laptop, your notepad and your smartphone, are Universal Turing Machines. And then the next step was to wonder if that machine can be said to be "intelligent", i.e. can behave like a human being (Turing's Test), can have conscious states, and can be even smarter than its creator (the Singularity).

The second motivation was purely business. Automation has been a source of productivity increase and wealth creation since ancient times. It accelerated during the industrial revolution and it still is an important factor in economic development. There isn't a day when a human being isn't replaced by a machine. Machines work 24 hours a day and 7 days a week, don't go on strike, don't

have to stop for lunch, don't have to sleep, don't get sick, don't get angry or sad. Either they function or they don't. If they don't, we simply replace them with other machines. Automation was pervasive in the textile industry way before computers were invented. Domestic appliances automated household chores. Assembly lines automated manufacturing. Agricultural machines automated grueling rural chores. That trend continues. As i type, machines (sensing cameras hanging from traffic lights remotely connected to the traffic division of a city) are replacing traffic police in many cities of the world to catch drivers who don't stop at red lights.

A third motivation was idealistic. An "expert system" could provide the service that the best expert in the world provides. The difference is that the human expert cannot be replicated all over the world, the expert system could. Imagine if we had an expert system that clones the greatest doctors in the world and then made that expert system available for free to the world's population, 24 hours a day, 7 days a week.

Brain Simulation and Intelligence

Behind the approach of neural networks is the hidden assumption that intelligence and perhaps consciousness itself arise out of complexity. This is a notion that dates back to at least William Grey-Walter who in 1949, before the age of digital computers, already designed early robots named *Machina Speculatrix* using analogue electronic circuit to simulate brain processes. In 2012 David Deamer calculated the "brain complexity" of several animals ("Consciousness and intelligence in mammals: Complexity thresholds").

All the "intelligent" brains that we know are made of neurons. Could the brain be made of ping-pong balls and still be as intelligent? If we connect a trillion ping-pong balls do we get a conscious being? Presumably not. What if the ping-pong balls are made of a material that conducts electricity? Presumably still not.

What if i connect them exactly like the neurons are connected in my brain: do i get a duplicate of my consciousness or at least a being that is as "intelligent" as me? The hidden assumption behind neural networks is that the material doesn't matter, that it doesn't have to be neurons (flesh), at least insofar as intelligence is concerned; hence, according to the connectionist model, that system of trillion of ping-pong balls would be as intelligent as me, as long as it duplicates exactly what happens in my brain.

The Body

A lot of what books on machine intelligence say is based on a brain-centered view of the human being. I may agree that my brain is the most important organ of my body (i'm ok with transplanting just about any organ of my body except my brain). However, this is probably not what evolution had in mind. The brain is one of the many organs designed to keep the body alive so that the body can find a mate and make children. The brain is not the goal but one of the tools to achieve that goal.

Focusing only on mental activities when comparing humans and machines is a categorical mistake. Humans do have a brain but don't belong to the category of brains: they belong to the category of animals, which are mainly recognizable by their bodies. Therefore, one should compare machines and humans based on bodily actions and not just on the basis of printouts, screenshots and files. Of all the things that i do during a day (from running to reading a book) what can a machine do? what will a machine be able to do in ten years? in 20 years? in 200 years? i suspect we are very far from the day that a machine can simply play soccer in any meaningful way with ten-year old children, let alone with champions. Playing a match of chess with the world champion of chess is actually easy. It is much harder for a machine to do any of the things that we routinely do in our home (that our bodies do).

Furthermore, there's the meaning of action. The children who play soccer actually enjoy it. They scream, they are competitive, they cry

if they lose, they can be mean, they can be violent. There is passion in what we do. Will an android that plays decent soccer in 3450 (that's a realistic date in my opinion) also have all of that? Let's take something simpler, that might happen in 50 or 100 years: at some point we'll have machines capable of reading a novel; but will they understand what they are reading? Is it the same "reading" that i do? This is not only a question about the self-awareness of the machine but about what the machine will do with the text it reads. I can find analogies with other texts, be inspired to write something myself, send the text to a friend, file it in a category that interests me. There is a follow-up to it. Machines that read a text and simply produce an abstract representation of its content (and we are very far from the day when they will be able to do so) are useful only for the human who will use it.

The same considerations apply to all the corporeal activities that are more than simple movements of limbs.

The body is the reason why i think the Turing Test is not very meaningful. The Turing Test locks a computer and a human being in two rooms, and, by doing so, it removes the body from the test. My test (let's immodestly call it the Scaruffi Test) would be different: we give a soccer ball to both the robot and the human and see who dribbles better. I am not terribly impressed that a computer beat the world champion of chess (i am more impressed with the human, that it took so long for a machine with virtually infinite memory and processing power to beat a human). I will be more impressed the day a robot dribbles better than Lionel Messi.

If you remove the body from the Turing test, you are removing pretty much everything that defines a human being as a human being. A brain kept in a jar is not a human being: it is a gruesome tool for classrooms of anatomy.

(I imagine my friends at the nearest A.I. lab already drawing sketches of robots capable of intercepting balls and then of kicking them with absolute precision towards the goal and with such force that no goalkeeper could stop them, but that's precisely what we don't normally call "intelligence", that is precisely what clocks and

photocopiers do, i.e. they can do things that humans cannot do such as keeping accurate time and making precise copies of documents, and that is not yet what Messi does when he dribbles defenders).

Intermezzo: We may Overestimate Brains

The record for brain size compared with body mass does not belong to Homo Sapiens: it belongs to the squirrel monkey (5% of the body weight, versus 2% for humans). The sparrow is a close second.

The longest living beings on the planet (bacteria and trees) have no brain.

Childhood

A machine is born adult. It's the equivalent of you being born at 25; and never aging until the day that an organ stops working. One of the fundamental facts of human intelligence is that it comes via childhood. First we are children, then we get the person that is writing (or reading) these lines. The physical development of the body goes hand in hand with the cognitive development of the mind. The developmental psychologist Alison Gopnik has written in "The Philosophical Baby" (2009) that the child's brain is wildly different from the adult brain (in particular the prefrontal cortex). She even says that they are two different types of Homo Sapiens, the child and the adult. They physically perform different functions. Whether it is possible to create "intelligence" equivalent to human intelligence without a formation period is debatable.

Alison Gopnik emphasized the way children learn about both the physical world and the social world via a process of "counterfactuals" (what ifs): they understand how the (physical and psychological) worlds function, then they create hypothetical ones (imaginary worlds and imaginary friends), then they are ready to create real ones (act in the world to change it and act on people to

change their minds). When we are children, we learn to act "intelligently" on both the world and on other people. Just like everything else with us, this function is not perfect. For example, one thing we learn to do is to lie: we lie in order to change the minds around us. A colleague once proudly told me: "Machines don't lie." That is one reason why I think that science is still so far away from creating intelligent machines. To lie is something you learn to do as a child, among many other things, among all the things that are our definition of "intelligence".

The State of the Art in Artificial Intelligence

Knowledge-based systems did not expand as expected: the human experts were not terribly excited at the idea of helping construct clones of themselves, and, in any case, the clones were not terribly reliable.

Expert systems also failed because of the World-wide Web: you don't need an expert system when thousands of experts post the answer to all possible questions. All you need is a good search engine. That search engine plus those millions of items of information posted (free of charge) by thousands of people around the world do the job that the "expert system" was supposed to do. The expert system was a highly intellectual exercise in representing knowledge and in reasoning heuristically. The search engine has no pretense of sophisticated logic but, thanks to the speed of today's computers and networks, it "will" find the answer. Within the world of computer programs it is a brute that can do the job once reserved to artists.

Note that the apparent "intelligence" of the Web (its ability to provide all sorts of questions) arises from the "non-intelligent" contributions of thousands of people in a way very similar to how the intelligence of an ant colony emerges from the non-intelligent contributions of thousands of ants.

In retrospect a lot of sophisticated logic-based software had to do with slow and expensive machines. As machines get cheaper and

faster and smaller, we don't need sophisticated logic anymore: we can just use fairly dumb techniques to achieve the same goals. As an analogy, imagine if cars, drivers and gasoline were very cheap and goods were provided for free by millions of people: it would be pointless to try and figure out the best way to deliver a good to a destination because one could simply ship many of those goods via many drivers with excellent chances that at least one good would be delivered on time at the right address. The route planning and the skilled knowledgeable driver would become useless, which is precisely what has happened in many fields of expertise in the consumer society: when is the last time you used a cobbler or a watch repairman?

The motivation to come up with creative ideas in A.I. was due to slow, big and expensive machines. Now that machines are fast, small and cheap the motivation to come up with creative ideas is much reduced. Now the real motivation is to have access to thousands of parallel processors and let them run for months. Creativity is now required in coordinating those processors to search through billions of items of information. The machine intelligence required in the world of cheap computers has become less of a logical intelligence and more of a logistical intelligence.

Meanwhile, in the 1980s some conceptual breakthroughs fueled real progress in robotics. Valentino Breitenberg, in his "Vehicles" (1984), showed that no intelligence is required for producing "intelligent" behavior: all that is needed is a set of sensors and actuators. As the complexity of the "vehicle" increases, the vehicle seems to display an increasingly intelligent behavior. Starting in about 1987, Rodney Brooks began to design robots that use little or no representation of the world. One can know nothing, and have absolutely no common sense, but still be able to do interesting things if equipped with the appropriate set of sensors and actuators.

The 1980s also witnessed a progressive rehabilitation of neural networks, a process that turned exponential in the 2000s. The discipline was rescued in 1982 by John Hopfield, who described a new generation of neural networks, based on simulating the

physical process of annealing. These neural networks were immune to Minsky's critique. Hopfield's key intuition was to note the similarity with statistical mechanics. Statistical mechanics translates the laws of Thermodynamics into statistical properties of large sets of particles. The fundamental tool of statistical mechanics (and soon of this new generation of neural networks) is the Boltzmann distribution (actually discovered by Josiah-Willard Gibbs in 1901), a method to calculate the probability that a physical system is in a specified state. Building on Hopfield's ideas, in 1983 Geoffrey Hinton and Terry Sejnowski developed the Boltzmann machine (technically, a Monte Carlo version of the Hopfield network), a software technique for networks capable of learning, and in 1986 Paul Smolensky introduced a further optimization, the Restricted Boltzmann Machine. These were carefully calibrated mathematical algorithms to build neural networks to be both feasible (given the dramatic processing requirements of neural network computation) and plausible (that solved the problem correctly).

This school of thought merged with another one that was coming from a background of statistics and neuroscience. Credit goes to Judea Pearl for introducing Bayesian thinking into Artificial Intelligence to deal with probabilistic knowledge ("Reverend Bayes on Inference Engines", 1982).

The Swedish statistician Ulf Grenander (who in 1972 had established the Brown University Pattern Theory Group) fostered a conceptual revolution in the way a computer should describe knowledge of the world: not as concepts but as patterns. His "general pattern theory" provided mathematical tools for identifying the hidden variables of a data set. Grenander's pupil David Mumford studied the visual cortex and came up with a hierarchy of modules in which inference is Bayesian (Thomas Bayes being the 18th century mathematician who developed Probability Theory as we know it today) and it is propagated both up and down ("On the computational architecture of the neocortex II", 1992). The assumption was that feedforward/feedback loops in the visual region integrate top-down expectations and bottom-up observations

via probabilistic inference. Basically, Mumford applied hierarchical Bayesian inference to model how the brain works. Hinton's Helmholtz machine of 1995 was de facto an implementation of those ideas: an unsupervised learning algorithm to discover the hidden structure of a set of data based on Mumford's and Grenander's ideas. The hierarchical Bayesian framework was later refined by Tai-Sing Lee of Carnegie Mellon University ("Hierarchical Bayesian inference in the visual cortex", 2003). These studies were also the basis for the widely-publicized "Hierarchical Temporal Memory" model of the startup Numenta, founded in 2005 in Silicon Valley by Jeff Hawkins, Dileep George and Donna Dubinsky. It was another path to get to the same paradigm: hierarchical Bayesian belief networks.

The field did not take off until 2006, when Geoffrey Hinton developed Deep Belief Networks, a fast learning algorithm for Restricted Boltzmann Machines. What had truly changed between the 1980s and the 2000s was the speed (and the price) of computers. Hinton's algorithms worked wonders when used on thousands of parallel processors. That's when the media started publicizing all sorts of machine learning feats.

Deep Belief Networks are layered hierarchical architectures that stack Restricted Boltzmann Machines one on top of the other, each one feeding its output as input to the one immediately higher, with the two top layers forming an associative memory. The features discovered by one RBM become the training data for the next one.

DBNs are still limited in one respect: they are "static classifiers", i.e. they operate at a fixed dimensionality. However, speech or images don't come in a fixed dimensionality, but in a (wildly) variable one. They require "sequence recognition", i.e. dynamic classifiers, that DBNs cannot provide. One method to expand DBNs to sequential patterns is to combine deep learning with a "shallow learning architecture" like the Hidden Markov Model.

Another thread in "deep learning" originated with Kunihiro Fukushima's convolutional networks ("Neocognitron - A Self-organizing Neural Network Model for a Mechanism of Pattern

Recognition Unaffected by Shift in Position", 1980) that led to Yann LeCun 's second generation convolutional neural networks ("Gradient-Based Learning Applied to Document Recognition", 1998). And Yeshua Bengio's stacked auto-encoders further improved the efficiency of deep learning ("Greedy Layer-wise Training of Deep Networks", 2007). In 2013 Google hired Hinton and Facebook hired LeCun.

Deep Belief Nets are probabilistic models that consist of multiple layers of probabilistic reasoning. Thomas Bayes' theorem of the 18th century is rapidly becoming one of the most influential scientific discoveries of all times (not bad for an unpublished manuscript discovered after Bayes' death). Bayes' theory of probability basically interprets knowledge as a set of probabilistic (not certain) statements and interprets learning as a process to refine those probabilities. As we acquire more evidence, we refine our beliefs. The developmental psychologist Jenny Saffran showed that babies use probability theory to learn about the world, and they do learn very quickly a lot of facts ("Statistical Learning By 8-Month-Old Infants", 1996). So Bayes had stumbled on an important fact about the way the brain works, not just a cute mathematical theory.

The story of robots is similar. Collapsing prices and increased speeds have enabled a generation of robots based on relatively old technology (in some cases really old technology, like the gyroscope, invented 150 years earlier). Cynthia Breazeal's emotional robot "Kismet" (2000), Ipke Wachsmuth's conversational agent "Max" (2004), Honda's humanoid robot "Asimo" (2005), Osamu Hasegawa's robot that learned functions it was not programmed to do (2011) and Rodney Brooks' hand programmable robot "Baxter" (2012) sound good on paper but look still as primitive as Shakey in person. Manufacturing plants have certainly progressed dramatically and can build, at a fraction of the cost, what used to be unfeasible; but there has been no conceptual breakthrough. All the robots built today could have been built in 1980 if the same manufacturing techniques had been available. There is little that is conceptually new (other than understanding

more about the smooth movements of the limbs). What is truly new is the techniques of advanced manufacturing and the speed of computers.

The field of "genetic algorithms", or, better, evolutionary computing, has witnessed progress that mirrors the progress in neural network algorithms; notably, in 2001 Nikolaus Hansen introduced the evolution strategy called "Covariance Matrix Adaptation" (CMA) for numerical optimization of non-linear problems. This has been widely applied to robotic applications and certainly helped better calibrate the movements of robots.

Despite all the hoopla, to me machines are still way less "intelligent" than most animals. Recent experiments with neural networks were hailed as sensational triumphs because a computer finally managed to recognize a cat (at least a few times) after being presented thousands of videos of cats. How long does it take a mouse to learn how a cat looks like? And that's despite the fact that computers use the fastest possible communication technology, whereas the neurons of a mouse's brain use hopelessly old-fashioned chemical signaling.

One of the very first applications of neural networks was to recognize numbers. Sixty years later the ATM of my bank still cannot recognize the amounts on many of the cheques that I deposit, but any human being can. Ray Kurzweil is often (incorrectly) credited with inventing "optical character recognition" (OCR), a technology that dates back to the 1950s (the first commercial OCR system was introduced by David Shepard's Intelligent Machines Research Corporation and became the basis for the Farrington Automatic Address Reading Machine delivered to the Post Office in 1953, and the term "OCR" itself was coined by IBM for its IBM 1418 product). Buy the most expensive OCR software and feed it the easiest possible case: a well-typed page from a book or magazine. It will probably make some mistakes that humans don't make, but, more interestingly, now slightly bend a corner of the page and try again: any human can still read the text,

but the most sophisticated OCR software on the market will go berserk.

For similar reasons we still don't have machines that can read cursive handwriting, despite the fact that devices with handwriting recognition features already appeared in the 1990s (GO's PenPoint, Apple's Newton). Most people don't even know that their tablet or smartphone has such a feature: it is so inaccurate that very few people ever use it. And, yet, humans (even not very intelligent ones) can usually read other people's handwriting with little or no effort.

As i write, Google's self-driving cars are the sensation in Silicon Valley, but this technology is at least 30 years old: Ernst Dickmanns demonstrated the robot car "VaMoRs" in 1986 and in october 1994 his modified Mercedes drove the Autoroute 1 near Paris in heavy traffic at speeds up to 130 km/h. (Incidentally, Google engineers still use the "miles" of the ancient imperial system instead of the kilometers of the metric system, a fact that hardly qualifies as "progress" to me).

We are flooded with news of robots performing all sorts of human tasks, except that most of those tasks are useless. On the other hand, commenting on the ongoing unmanned Mars mission, in April 2013 NASA planetary scientist Chris McKay told me that "what Curiosity has done in 200 days a human field researcher could do in an easy afternoon." And that is the most advanced robotic explorer ever built.

We also have to realize that what looks astonishing to us might not look so impressive to future generations. Today very few people are amazed by cars. When the car was invented, instead, it must have been an incredible sight: an object that moves with no animal pulling it and no peasants pushing it. How could that be possible? For thousands of years people had seen objects move only when they were attached to ropes pulled by humans or animals. But now an object was moving by itself on four wheels. The people who today marvel at the sight of a self-driving car are behaving like the first people who were shocked at the sight of a self-moving object.

Today, like back then, the temptation is to call it "magic" (or "Artificial Intelligence"). In reality, it is just an application of well-known physical laws.

It is interesting how different generations react to the stupidity of machines: the generation that grew up without electronic machines gets extremely upset (because the automated system can be so much more stupid than the old-fashioned manual system), my generation (that grew up with machines) gets somewhat upset (because machines are not much smarter than they were when i was a kid), and the younger generations are progressively less upset, with the youngest ones simply taking for granted that customer support has to be what it is (from lousy to non-existent) and that many things (pretty much all the things that require common sense, expertise, and what we normally call "intelligence") are virtually impossible for machines. Hence the different attitudes towards the fact that the most widespread computing platform of 2012, Facebook, does not allow you to edit your posts and comments (you need to delete them and re-post them, like in the age of punched cards). One could write a chapter on "The State of Machine Stupidity" instead of "The State of Machine Intelligence" and it might be a longer one.

What the new A.I. does is very simple: lots of number crunching. It is a smart way to manipulate large datasets for the purpose of classification. It was not enabled by a groundbreaking paradigm shift but simply by increased computing power. The "Google Brain" project started at Google by Andrew Ng in 2011 is the quintessential example of this approach. In June 2012 a combined Google/Stanford research team used an array of 16,000 processors to create a neural network with more than one billion connections and let it loose on the Internet to learn from millions of YouTube videos how to recognize cats. Given the cost, size and speed of computers back then, 30 years ago nobody would have tried to build such a system. The difference between then and now is that today A.I. scientists can use thousands of powerful computers to get what they want. It is, ultimately, brute force with little or no

sophistication. Whether this is how the human mind does it is debatable. And, again, we should be impressed that 16,000 of the fastest computers in the world took a few months to recognize a cat, something that a kitten with a still undeveloped brain can do in a split second. I would be happy if the 16,000 computers could just simulate the 302-neuron brain of the roundworm, no more than 5000 synapses that nonetheless can recognize with incredible accuracy a lot of very interesting things.

Brute force is the paradigm that now dominates A.I. After all, by indexing millions of webpages, a search engine is capable of providing an answer to the vast majority of questions (even "how to" questions), something that no expert system came close to achieving.

One wonders if slow and cumbersome computers were a gift to the scientific community of the 1960s because those archaic machines forced computer scientists to come up with creative models instead of just letting modern high-speed computers crunch numbers until a solution is found.

Most disheartening has been the insignificant progress in Natural Language Processing (NLP) since 1970.

Following are examples of exchanges between a person and Terry Winograd's program SHRDLU in 1970:

Person: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.

Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.

Person: DOES THE SHORTEST THING THE TALLEST PYRAMID'S SUPPORT SUPPORTS SUPPORT ANYTHING GREEN?

Computer: YES, THE GREEN PYRAMID.

Person: WHAT COLOR IS IT?

Computer: BY "IT", I ASSUME YOU MEAN THE SHORTEST THING THE TALLEST PYRAMID'S SUPPORT SUPPORTS. RED

Person: DOES A STEEPLE

Computer interrupting: SORRY, I DON'T KNOW THE WORD "STEEPLE".

Person: A "STEEPLE" IS A STACK WHICH CONTAINS TWO GREEN CUBES AND A PYRAMID.

Computer: i UNDERSTAND.

Person: ARE THERE ANY STEEPLES NOW?

Computer: NO.

It beats many of today's "personal digital assistants".

We virtually abandoned the idea of having machines understand and speak our language and resorted to the opposite: make humans speak like machines. That is what you do when you talk on the phone with a machine that asks you for numbers and that's what you do when you talk to your smartphone's "assistant" according to the rules of that assistant. Nothing illustrates the gap between machine intelligence and human intelligence better than comparing how much of a language a toddler learns in two years and how little of any language all machines ever invented have learned in more than 60 years.

IBM's Watson, that debuted with much fanfare in 2011 on a quiz show competing against human experts, was actually not capable of understanding the spoken questions: the questions were delivered to Watson as text files, not as spoken questions (which, of course, distorted the whole game).

The most popular search engines are still keyword-based. Progress in search engines has been mainly in indexing and ranking webpages, not in understanding what the user is looking for. Try for example "Hey i had a discussion with a friend about whether Qaddafi wanted to get rid of the US dollar and he was killed because of that" and see what you get (as i write these words, Google returns first of all my own website with the exact words of that sentence and then a series of pages that discuss the assassination of the US ambassador in Libya) . Communicating with a search engine is a far (far) cry from communicating with human beings.

Announcements of products that can understand natural language, such as SIRI for Apple's iPhone, have been wild

exaggerations. These products understand only the most elementary of sounds, and only sometimes, just like their ancestors of decades ago. Promising that a device will be able to translate speech on the fly (like Samsung did with its Galaxy S4 in 2013) is a good way to embarrass yourself and to lose credibility among your customers.

To be fair, progress in natural language understanding was hampered by the simple fact that humans prefer not to speak to another human in our time-consuming natural language. Sometimes we prefer to skip the "Good morning, how are you?" and get straight to the "Reset my modem" in which case saying "One" to a machine is much more effective than having to wait for a human operator to pick up the phone and to understand your issue. Does anyone actually understand the unintelligible announcements in the New York subway? Communicating in natural language is not always a solution, as SIRI users are rapidly finding out on their smartphone. Like it or not, humans can more effectively go about their business using the language of machines. For a long time, therefore, Natural Language Processing remained an underfunded research project with few visible applications. It was only recently that interest in "virtual personal assistants" has resurrected the field.

Machine Translation too has disappointed. Despite recurring investments in the field by major companies, your favorite online translation system succeeds only with the simplest sentences, just like Systran in the 1970s. Here are some random Italian sentences from my old books translated into English by the most popular translation engine: "Graham Nash the content of which led nasal harmony", "On that album historian who gave the blues revival", "Started with a pompous hype on wave of hippie phenomenon".

The method that has indeed improved the quality of automatic translation is the statistical one, pioneered in the 1980s by Fred Jelinek's team at IBM ("A statistical approach to language translation", 1988). When there are plenty of examples of (human-made) translations, the computer can perform a simple statistical analysis and pick the most likely translation. Note that the computer

isn't even trying to understand the sentence: it has no clue whether this is about languages or parliamentary elections. The statistical approach works wonders when there are thousands of (human-made) translations of a sentence, for example between Italian and English. It works awfully when there are fewer, like in the case of Chinese to English.

Even if we ever get to the point that a machine can translate a sentence, here is the real test: "'Thou' is an ancient English word". Translate that into Italian as "'Tu' e` un'antica parola Inglese" and you get an obviously false statement ("Tu" is not an English word). The trick is to understand what the original sentence means, not to just mechanically replace English words with Italian words. If you understand what it means, then you'll translate it as "'Thou' e` un'antica parola Inglese", i.e. you don't translate the "thou"; or, depending on the context, you might want to replace "thou" with an ancient Italian word like "'Ei' e` un'antica parola Italiana" (where "ei" actually means "he" but it plays a similar role to "thou" in the context of words that changed over the centuries). A machine will be able to get it right only when it fully understands the meaning and the purpose of the sentence, not just its structure.

(There is certainly at least one quality-assurance engineer who, informed of this passage in this book, will immediately enter a subroutine in the machine translation program to correctly translate "'Thou' is an ancient English word". That is precisely the brute-force approach that I was talking about).

Or take Ronald Reagan's famous sarcastic statement, that the nine most terrifying words in the English language are "I'm from the government and I'm here to help". Translate this into Italian and you get "Le nove parole piu` terrificanti in Inglese sono `io lavoro per il governo e sono qui per aiutare". Those are neither nine in the Italian translation (they are ten) and they are not "Inglese" (English) because they are now Italian. An appropriate translation would be "Le dieci parole piu` terrificanti in Italiano sono `io lavoro per il governo e sono qui per aiutare". Otherwise the translation, while technically impeccable, makes no practical sense.

Or take Bertrand Russell's paradox: "the smallest positive integer number that cannot be described in fewer than fifteen words". This is a paradox because the sentence in quotes contains fourteen words. Therefore if such an integer number exists, it can be described by that sentence, which is fourteen words long. When you translate this paradox into Italian, you can't just translate fifteen with "quindici". You first need to count the number of words. The literal translation "il numero intero positivo piu` piccolo che non si possa descrivere in meno di quindici parole" does not state the same paradox because this Italian sentence contains sixteen words, not fourteen like the original English sentence. You need to understand the meaning of the sentence and then the nature of the paradox in order to produce an appropriate translation. I could continue with more and more convoluted self-referential sentences that can lead to trivial mistakes when translated "mechanically" without understanding what they are meant to do.

To paraphrase the physicist Max Tegmark, a good explanation is one that answers more than was asked. If i ask you "Do you know what time it is", a "Yes" is not a good answer. I expect you to at least tell me what time it is, even if it was not specifically asked. Better: if you know that i am in a hurry to catch a train, i expect you to calculate the odds of making it to the station in time and to tell me "It's too late, you won't make it" or "Run!" If i ask you "Where is the library?" and you know that the library is closed, i expect you to reply with not only the location but also the important information that it is currently closed (it is pointless to go there). If i ask you "How do i get to 330 Hayes St?" and you know that it used to be the location of a popular Indian restaurant that just shut down, i expect you to reply with a question "Are you looking for the Indian restaurant?" and not with a simple "It's that way". If i am in a foreign country and ask a simple question about buses or trains, i might get a lengthy lecture about how public transportation works. Speaking a language is pointless if one doesn't understand what language is all about. A machine can easily be programmed to answer the question "Do you know what time it is" with the time (and not a

simple "Yes"), and it can easily be programmed to answer similar questions with meaningful information; but we "consistently" do this for all questions, and not because someone told us to answer the former question with the time and other questions with meaningful information, but because that is what our intelligence does: we use our knowledge and common sense to formulate the answer. In the near future it will be extremely difficult to build machines that can understand the simplest of sentences. It may take centuries before we have a machine that can have a conversation like the ones I have with my friends on the Singularity. And that would still be a far cry from what humans do: consistently provide an explanation that answers more than it was asked.

A lot more is involved than simply understanding a language. If people around me speak Chinese, they are not speaking to me. But if one says "Sir?" in English, and i am the only English speaker around, i am probably supposed to pay attention.

The state of Natural Language Processing is will be represented by the results returned by the most advanced search engines: the vast majority of results are precisely the kind of commercial pages that i don't want to see. Which human would normally answer "do you want to buy perfume Katmandu" when i inquire about Katmandu's monuments? It is virtually impossible to find out which cities are connected by air to a given airport because the search engines all simply return hundreds of pages that offer "cheap" tickets to that airport.

I like to discuss with machine-intelligence fans a simple situation. Let's say you are accused of a murder you did not commit. How many years will it take before you are willing to accept a jury of 12 robots instead of 12 humans? Initially, this sounds like a question about "when will you trust robots to decide whether you are guilty or innocent?" but it actually isn't (i would probably trust a robot better than many of the jurors who are easily swayed by good looks, racial prejudices and many other unpredictable factors). The question is about understanding the infinite subtleties of legal debates, the language of lawyers and, of course, the language of the witnesses.

The odds that those 12 robots fully understand what is going on at a trial will remain close to zero for a long time.

In my opinion the "footnotes" in the history of Artificial Intelligence were not just footnotes: they were colossal failures. They were all great ideas. In fact, they were probably the "right" ideas: of course, an intelligent machine must be capable of conversing in natural language; of course, it must be able to walk around, look for food, and protect itself; of course, it must be able to understand what people say (each person having a slightly different voice); of course, it would make more sense for software to "evolve" by itself than to be written by someone (just like any form of intelligent life did); of course, we would expect an intelligent machine to be able to write software (and build other machines, like we do); of course, it would be nice if the machine were capable of translating from one language to another; of course, it would make sense to build a computer that is a replica of a human brain if what we expect is a performance identical to the performance of a human brain.

These ideas remained footnotes for a simple reason: from the point of view of scientific ideas, the last three decades of Artificial Intelligence research have been a failure.

"Machines will be capable, within twenty years, of doing any work that a man can do" (Herbert Simon, 1965). Slightly optimistic back then. I haven't seen anything yet that makes me think that statement (by a Nobel Prize winner) is any more accurate today.

Note that, ironically, it was A.I. that made computers popular and fueled progress in computer science. The idea of a thinking machine, not the usefulness of computers, drove the initial development. Since those days progress in A.I. has been scant, but computers have become household appliances. Your notepad and smartphone are accidental byproducts of a failed scientific experiment.

An Easy Science

When a physicist makes a claim, an entire community of physicists is out there to check that claim. The paper gets published only if it survives peer review, and usually many months after it was written. A discovery is usually accepted only if the experiment can be repeated elsewhere. For example, when OPERA announced particles traveling faster than light, the whole world conspired to disprove them, and eventually it succeeded. It took months of results before CERN accepted that probably (not certainly) the Higgs boson exists.

Artificial Intelligence practitioners, instead, have a much easier life. Whenever they announce a new achievement, it is largely taken at face value by the media and by the A.I. community at large. If a computer scientist announces that her or his program can recognize a cat, the whole world posts enthusiastic headlines even if nobody has actually seen it in action, and nobody has been able to measure and doublecheck its performance: can i make a video of a cat and feed it into the program and see if it recognizes the cat? When in 2012 Google announced that "Our vehicles have now completed more than 300,000 miles of testing" (a mile being 1.6 kilometers for the rest of the world), the media simply propagated the headline without asking simple questions such as "starting when?" or "under which conditions?" or "on which roads?" "at what time of the day"? Most people now believe that self-driving cars are feasible even though they have never been in one. Many of the same people probably don't believe all the weird consequences of Relativity and Quantum Mechanics, despite the many experiments that confirmed them.

The 2004 DARPA challenge for driverless cars was staged in the desert between Los Angeles and Las Vegas (i.e. with no traffic). The 2007 DARPA urban challenge took place at the George Air Force Base. I'll let you search videos on the Web and let you decide if what you see looks impressive. Interestingly, two highly educated friends told me that a DARPA challenge took place downtown Los Angeles in heavy traffic. That never took place. Too often the belief in the feats of A.I. systems feels like the stories of devout people

who saw an apparition of a saint and all the evidence you can get is a blurred photo.

In 2005 the media reported that Hod Lipson at Cornell University had unveiled the first "self-assembling machine" (the same scientist in 2007 also unveiled the first "self-aware" robot), and in 2013 the media reported that the "M-blocks" developed at the MIT by John Romanishin, Kyle Gilpin and Daniela Rus were self-constructing machines. Unfortunately, these reports were wild exaggerations.

In May 1997 the IBM supercomputer "Deep Blue" beat then chess world champion Garry Kasparov in a widely publicized match. What was less publicized is that the match was hardly fair: Deep Blue had been equipped with an enormous amount of information about Kasparov's chess playing, whereas Kasparov knew absolutely nothing of Deep Blue; and during the match IBM engineers kept tweaking Deep Blue with heuristics about Kasparov's moves. Even less publicized were the rematches, in which the IBM programmers were explicitly forbidden to modify the machine in between games. The new more powerful versions of Deep Blue (renamed Frintz) could beat neither Vladimir Kramnik, the new world chess champion, in 2002 nor Kasparov himself in 2003. Both matches ended in a draw. What is incredible to me is that a machine equipped with virtually infinite knowledge of the game and of its opponent, and with lightning-speed circuits that can process virtually infinite number of moves in a split second cannot beat a much more rudimentary object such as the human brain equipped with a very limited and unreliable memory: what does it take for a machine to outperform humans despite all the technological advantages it has? Divine intervention? Nonetheless, virtually nobody in the scientific community (let alone in the mainstream media) questioned the claim that a machine had beaten the greatest chess player in the world.

If IBM is correct and, as it claimed at the time, Deep Blue could calculate 200 million positions per second whereas Kasparov's brain could only calculate three per second, who is smarter, the one who can become the world's champion with just three calculations

per second or the one who needs 200 million calculations per second? What Deep Blue certainly achieve was to get better at chess than its creators. But that is true of the medieval clock too, capable of keeping the time in a way that no human brain could, and of many other tools and machines.

Finding the most promising move in a game of chess is a lot easier than predicting the score of a Real Madrid vs Barcelona game, something that neither machines nor humans are even remotely close to achieving. The brute force of the fastest computers is enough to win a chess game, but the brute force of the fastest computers is not enough to get a better soccer prediction than, say, the prediction made by a drunk soccer fan in a pub. Ultimately what we are contemplating when a computer beats a chess master is still what amazed the public of the 1950s: the computer's ability to run many calculations at lightning speed, something that no human being can do.

IBM's Watson of 2013 consumes 85,000 Watts compared with the human brain's 20 Watts. For the televised match with the human experts Watson was equipped with 200 million pages of information including the whole of Wikipedia; and, in order to be fast, all that knowledge had to be stored on RAM, not on disk storage. The human experts who competed against Watson did not have access to all that information. De facto the human experts were not playing against one machine but against a whole army of machines, enough machines working to master and process all those data. A fairer match would be to pit Watson against thousands of human experts, chosen so as to have the same amount of data. And, again, the questions were conveniently provided to the machine as text files instead of spoken language. If you use the verb "to understand" the way we normally use it, Watson never understood a single question. And those were the easiest possible questions, designed specifically to be brief and unambiguous (unlike the many ambiguities hidden in ordinary human language). Watson didn't even hear the questions (they were written to it), let alone understand what the questioner was asking.

Over the decades i have personally witnessed several demos of A.I. systems that required the audience to simply watch and listen: only the creator was allowed to operate the system.

Furthermore, some of the most headline-capturing Artificial Intelligence research is supported by philanthropists at private institutions with little or no oversight by academia.

Personally, i wouldn't even know how to validate an announcement like "this system can recognize a cat": what percentage? and percentage of what? Can we perform the Turing test of cat recognizing? Can we compare who is best at recognizing cats, a chimp or the program? a mouse or the program? a fly or the program?

Many of the A.I. systems of the past have never been used outside the lab that created them. Their use by the industry, in particular, has been virtually nil.

Very few people bother to doublecheck the claims of the robotics and A.I. community. The media have a vested interest that the story be told (it sells) and the community as a whole has a vested interest that government and donors believe in the discipline's progress so that more funds will be poured into it.

Paul Nunez in "Brain, Mind, and the Structure of Reality" (2010) distinguishes between Type 1 scientific experiments and Type 2 experiments. Type 1 is an experiment that has been repeated at different locations by different teams and still holds. Type 2 is an experiment that has yielded conflicting results at different laboratories. UFO sightings, levitation tales and exorcisms are not scientific, but many people believe in their claims, and i will call them "Type 3" experiments, experiments that cannot be repeated by other scientists. Most of Artificial Intelligence experiments fall between Type 2 and Type 3.

News of the feats achieved by machines rapidly propagate worldwide thanks to enthusiastic bloggers and tweeters the same way that news about telepathy and levitation used to spread rapidly worldwide thanks to word of mouth without the slightest requirement of proof. (There are still millions of people who believe that cases of

levitation have been documented even though there is no footage and no witness to be found anywhere). The belief in miracles worked the same way: people wanted to believe that a saint had performed a miracle and they transmitted the news to all their acquaintances in a state of delirious fervor without bothering to doublecheck the facts and without providing any means to doublecheck the facts (address? date? who was there? what exactly happened?). The Internet is a much more powerful tool than the old "word of mouth" system. In fact, i believe that part of this discussion about machine intelligence is a discussion not about technology but about the World-wide Web as the most powerful tool ever invented to spread myths. And part of this discussion about machine intelligence is a discussion about the fact that 21st century humans want to believe that super-intelligent machines are coming the same way that people of previous centuries wanted to believe that magicians existed. The number of people whose infirmity has been healed after a visit to the sanctuary of Lourdes is very small (and in all cases one can find a simple medical explanation) but thousands of highly educated people still visit it when they get sick, poor or depressed. On 13 October 1917 in Fatima (Portugal) tens of thousands of people assembled because three shepherd children had been told by the Virgin Mary (the mother of Jesus) that she would appear at high noon. Nobody saw anything special (other than the Sun coming out after a rain) but the word that a miracle had taken place in Fatima spread worldwide. Believe it or not, that is pretty much what happens in 2013 when a fanatical blogger reports a feat performed by an A.I. software or a robot as a new step towards the Singularity. The people like me who remain skeptical of the news are looked down upon the same way that skeptics were looked down upon after Fatima: "What? You still don't believe that the Virgin Mary appeared to those children? What is wrong with you?". Which, of course, shifts the burden of proof on the skeptic who is asked to explain why one would NOT believe in the miracle (sorry, i meant "in the machine's intelligence") instead of pressing the inventor/scientist/lab/firm into proving that the

miracle/feat has truly been accomplished and can be repeated at will and that it really does what bloggers said it does.

"Whenever a new science achieves its first big successes, its enthusiastic acolytes always fancy that all questions are now soluble" (Gilbert Ryle, "The Concept of Mind", 1949, six years before Artificial Intelligence was born).

Intermezzo and Trivia: the Original App

At the same time the real achievements of the machine are sometime neglected. I am not terribly impressed that computers can play chess. I am much more impressed that computers can forecast the weather, since the atmosphere is a much more complex system than the game of chess. The media have always devoted more attention to the game of chess because its rules are easier to explain to the general public, whereas the rules that guide air flow and turbulence are rather exotic. However, it turns out that weather forecasting was the original "app".

Weather forecast was the "mission impossible" of the early computers. The first weather forecast using a computer simulation dates back to March 1950, to the early history of electronic computers. The computer was an ENIAC and it took just about 24 hours to calculate the weather forecast for the next 24 hours. Weather forecasting was a particularly interesting application of electronic computing for John Von Neumann. In fact, it was "the" application originally envisioned for the machine that Von Neumann designed at Princeton's Institute for Advanced Studies (IAS), the machine that introduced the "Von Neumann architecture" still used today. Mathematicians had known for a while that solving this problem, i.e. modeling the air flow, required solving a non-linear system of partial differential equations - Lewis Richardson had

published the milestone study in this field, "Weather Prediction by Numerical Process" in 1922 - and that is why mathematicians thought this was an avantgarde problem; and that's why Von Neumann felt that solving it with a computer would not only help the community of meteorologists but also prove that the electronic computer was no toy. The ENIAC program, however, used an approximation devised by Jule Charney in 1948 ("On the scale of atmospheric motions"). A computer model for the general circulation of the atmosphere had to wait until 1955, when Norman Phillips, also at Princeton, presented his equations at the Royal Meteorological Society, and fed them into the IAS computer Maniac i ("The general circulation of the atmosphere", 1955). Meanwhile, the ability to predict the weather was dramatically improved in 1957 when the first satellite was launched. By 1963 a Japanese scientist at UCLA, Akio Arakawa, had tweaked Phillips' equations and written a Fortran program on an IBM 709, with help from IBM's Large Scale Scientific Computation Department in San Jose ("Computational Design for Long-Term Numerical Integration of the Equations of Fluid Motion", 1966). IBM was obviously ecstatic that their computer could be used to solve such a strategic problem as predicting the weather. It was the Fortran programming language's baptism of fire, as the 709 was the first commercial computer equipped with a Fortran compiler. At this point it looked like it was just a matter of waiting for computers to get faster. Alas, in the same year that Arakawa produced the first meaningful weather forecast, Edward Lorenz proved that the atmosphere belongs to the class of system now known as "chaotic" ("Deterministic Nonperiodic Flow", 1963): there is a limit to how accurately one can predict their behavior. In fact, as computers grow exponentially faster due to Moore's law, weather forecast models have not become exponentially more accurate. Robin Stewart has shown that "despite this exponential increase in computational power, the accuracy of forecasts has increased in a decidedly linear fashion" ("Weather Forecasting by Computers", 2003). Even today

meteorologists can only give us useful forecasts of up to about a week.

Note that, unlike chess and machine translation, this problem is not currently solved by using statistical analysis. It is solved by observing the current conditions and applying physical laws (as derived by those pioneering scientists). Statistical analysis requires an adequate sample of data, and a relatively linear behavior. Weather conditions, instead, are never the same, and the nonlinear nature of chaotic systems like the atmosphere makes it very easy to come up with grotesquely wrong predictions. This does not mean that it is impossible to predict the weather using statistical analysis; just that it is only one method out of many, a method that has been particularly successful in those fields where statistical analysis makes sense but was not feasible before the introduction of powerful computers. There is nothing magical about its success, just like there is nothing magical about our success in predicting the weather. Both are based on good old-fashioned techniques of computational math).

Don't be Fooled by the Robot

The bar is being set very low for robotics too. Basically, any remote-controlled toy (as intelligent as the miniature trains that were popular in the 1960s) is now being hailed as a step toward superhuman intelligence. I always advise robotics fans to visit the Musee Mecanique in San Francisco, that has a splendid collection of antique coin-operated automatic mechanical musical instruments... sorry, i meant "of robotic musicians", before they venture into a discussion about progress in robotics. They don't constitute what we normally call "intelligence".

Does driving a car qualify as a sign of "intelligence"? Maybe it does, but it has to be "really" what it means for humans. There is no car that has driven even one meter without help from humans. The real world is a world in which first you open the garage door, then you stop to pick up the newspaper, then you enter the street and

you will stop if you see a pedestrian waiting to cross the street. No car has achieved this skill yet. They self-drive only in highly favorable conditions on well marked roads with well marked lanes and only on roads that the manufacturing company has mapped accurately. And i will let you imagine what happens if the battery dies or there's a software bug... What does the self-driving car do if it is about to enter a bridge when an earthquake causes the bridge to collapse? Presumably it will just drive on. What does the self-driving car do if it is stopping at a red light and a man with a gun breaks the window? Probably nothing: it's a red light. If you fall asleep in a self-driving car, your chances of dying skyrocket. There are countless rules of thumb that a human driver employs all the time, and they are based on understanding what is going on. A set of sensors wrapped in a car's body is not understanding anything about what is going on. So even if driving a car qualified as an "intelligent" skill, machines have not achieved that skill yet.

Human-looking automata that mimic human behavior have been built since ancient times and some of them could perform sophisticated movements. They were mechanical. Today we have electromechanical sophisticated toys that can do all sort of things. There is a (miniature) toy that looks like a robot riding a bicycle. Technically speaking, the whole toy is the "robot". Philosophically speaking, there is no robot riding a bicycle. The robot-like thing on top of the bicycle is redundant, it is there just for show: you can remove the android and put the same gears in the bicycle seat or in the bicycle pedals and the bike with no passenger would go around and balance itself the exact same way: the thing that rides the bicycle is not the thing on top of the bike (designed to trick the human eye) but the gear that can be placed anywhere on the bike. The toy is one piece: instead of one robot, you could put ten robots on top of each other, or no robot at all. Any modern toy store has toys that behave like robots doing some amazing thing (amazing for a robot, ordinary for a human). It doesn't require intelligence: just Japanese or Swiss engineering. This bike-riding toy never falls, even when it is not moving. It is designed to always stand vertical.

Or, better, it falls when it runs out of battery. That's very old technology. If that's what we mean by "intelligent machines", then they have been around for a long time. We even have a machine that flies in the sky using that technology. Does that toy represent a quantum leap forward in intelligence? Of course, no. It is remotely controlled just like a television set. It never "learned" how to bike. It was designed to bike. And that's the only thing it can do. Ever. The only thing that is truly amazing in these toys is the miniaturization, not the "intelligence".

If you want this toy to do something else, you'll have to add more gears of a different kind, specialized in doing that other thing. Maybe it is possible (using existing technology or even very old mechanical technology) to build radio-controlled automata that have one million different gears to do every single thing that humans do and that all fit in a size comparable to my body's size. It would still be a toy.

A human being is NOT a toy (yet).

A Consumer's Rant Against the Stupidity of Machines: Reverting Evolution?

When you buy an appliance and it turns out that you have to do something weird in order to make it work, it is natural to dismiss it as "a piece of garbage". However, when it is something about computers and networks, you are supposed to stand in awe and respectfully listen to (or read) a lengthy explanation of what you are supposed to do in order to please the machine, which is usually something utterly convoluted bordering on the incredible. This double standard creates the illusion that machines are becoming incredibly smart when in fact mostly we are simply witnessing poor quality assurance (due to the frantic product lifecycles of our times) and often incredibly dumb design. You never know what is going to happen to your favorite application when you download an "update". New releases (which you are forced to adopt even if you are perfectly happy with the old release) often result in lengthy

detours trying to figure out how to do things that were trivial in the previous release (and that have been complicated by the software manufacturer for no other reason than to have a new release). A few week ago my computer displayed the message "Updating Skype... Just a moment, we're improving your Skype experience". How in heaven do they know that this will improve my Skype experience? Of course they don't. The reason they want you to move to a new release is different: it will certainly improve THEIR experience. Whether it will also improve mine and yours is a secondary issue. At the least, any change in the user interface will make it more difficult to do the things to which you were accustomed.

We live in an age in which installing a wireless modem can take a whole day and external hard disks get corrupted after a few months "if you use them too often" (as an associate told me at Silicon Valley's most popular computer store).

Reality check: here is the transcript of a conversation with Comcast's automated customer support (recorded in 2012).

"If you are currently a Comcast customer, press 1" [I press 1]

"Please enter the ten-digit phone number associated with your account" [I enter my phone number]

"OK Please wait just a moment while i access your account"

"For technical help press 1"

"For billing press 2" [I press 2]

If you are calling regarding important information about Xfinity etc press 1 [I press 2]

"For payments press 1"

"For balance information press 2"

"For payment locations press 3"

"For all other billing questions press 4" [I press 4]

"For questions about your first bill press 1"

"For other billing questions press 3" [I press 3]

"Thank you for calling Comcast. Our office is currently closed."

(You can listen to it at <https://soundcloud.com/scaruffi/comcast-customer-support>)

Note: Comcast was testing one of the best automated customer support systems in the world.

Based on the evidence, it is easier to believe that we still are in the stone age of computer science than to believe that we are about to witness the advent of superhuman intelligence in machines.

Incidentally, there are very important fields in which we haven't even achieved the first step of automation: getting rid of paper. Health care, for example, still depends on paper: your medical records are probably stored in old fashioned files, not the files made of zeros and ones but the ones made of cardboard. We are bombarded daily by news of amazing medical devices and applications that will change the way diseases are prevented, identified and healed, but for the time being we have seen very little progress in simply turning all those paper files into computer files that the patient can access from a regular computer or smartphone and then save, print, email or delete at will.

What we have done so far, and only in some fields, is to replace bits and pieces of human intelligence with rather unintelligent machines that can only understand very simple commands (less than what a two-year old toddler can understand) and can perform very simple tasks (mostly, additions).

In the process we are also achieving lower and lower forms of human intelligence, addicted to having technology simplify all sorts of tasks (more about this later). Of course, many people claim the opposite: from the point of view of a lower intelligence, what unintelligent machines do might appear intelligent.

The Singularity as the Outcome of Exponential Progress

Nonetheless, the Singularity gurus are driven to enthusiastic prognostications about the evolution of machines: machines will soon become intelligent and next they will become intelligent in a superhuman way, acquiring a form of intelligence so high that humans won't be able to understand. There is an obvious

disconnect between the state of the art and what the Singularity crowd predicts. We are not even remotely close to a machine that can troubleshoot and fix an electrical outage or simply your washing machine, let alone a software bug. We are not even remotely close to a machine that can operate any of today's complex systems without human supervision. One of the premises of the theory of the Singularity is that machines will not only become intelligent but will even build other, smarter machines by themselves; but right now we don't even have software that can write other software.

The jobs that have been automated are repetitive and trivial. And in most cases the automation of those jobs has required the user/customer to accept a lower (not higher) quality of service. Witness how customer support is rapidly being reduced to a "good luck with your product" kind of service. The more automation around you, the more you (you) are forced to behave like a machine to interact with machines, precisely because they are still so dumb. The reason that we have a lot of automation is that (in expensive countries like the USA, Japan and the European countries) it saves money: machines are cheaper than humans. Wherever the opposite is true, there are no machines. The reason we are moving to online education is not that university professors failed to educate their students but that universities are too expensive. And so forth: in most cases it is the business plan, not the intelligence of machines, that drives automation.

Wildly optimistic predictions are based on the exponential progress in the speed and miniaturization of computers. In 1965 Gordon Moore predicted that the processing power of computers would double every 18 months ("Moore's law"), and so far his prediction has been correct. Look closer and there is little in what they say that has to do with software. It is mostly a hardware argument. And that is not surprising: predictions about the future of computers have been astronomically wrong in both directions but, in general, the ones that were too conservative were about hardware, the ones that were too optimistic were about software. What is amazing about today's smartphones is not that they can do

what computers of the 1960s could not do (they can do pretty much the same kinds of things) but that they are small, cheap and fast. The fact that there are many more software applications downloadable for a few cents means that many more people can use them, a fact that has huge sociological consequences; but it does not mean that a conceptual breakthrough has been reached in software technology. It is hard to name one software program that exists today and could not have been written in Fortran fifty years ago. If it wasn't written, the reason is that it would have been too expensive or that some required hardware did not exist yet.

Accelerating technological progress has largely been driven by accelerating cost of labor, not by real scientific innovation. The higher labor costs go, the stronger the motivation to develop "smarter" machines. Those machines, and the underlying technologies, were already feasible ten or twenty or even thirty years ago, but back then it didn't make economic sense for them to be adopted.

There has certainly been a lot of progress in computers getting faster, smaller and cheaper; and this might continue for a while. Even assuming that this will continue "exponentially" (as the Singularity crowd is quick to claim), the argument that this kind of (hardware) progress is enough to make a shocking difference in terms of machine intelligence is based on an indirect assumption: that faster/smaller/cheaper will lead first to a human-level intelligence and then to a superior intelligence. After all, if you join together many many many dumb neurons you get the very intelligent brain of Albert Einstein. If one puts together millions of smartphones, maybe one gets superhuman intelligence. Maybe.

Incidentally, Moore's Law ("Cramming more components into integrated circuits", 1965) was widely predicted to continue in the foreseeable future, but its future does not look so promising anymore. It is not only that technology limits might be approaching. The original spirit behind Moore's Law was to show that the "cost" of making transistor-based devices would continue to decline. Even if the industry finds a way to continue to double the number of

transistors etched on a chip, the cost of doing so might start increasing soon: the technologies to deal with microscopic transistors are inherently expensive.

The discussion about the Singularity is predicated upon the premise that machines will soon be able to perform "cognitive" tasks that were previously exclusive to humans. This, however, has already happened. We just got used to it. The early computers of the 1950s were capable of computations that traditionally only the smartest and fastest mathematicians could even think of tackling, and the computers quickly became millions of times faster than the fastest mathematician. If computing is not an "exclusively human cognitive task", i don't know what would qualify. Since then computers have been programmed to perform many more of the tasks that used to be exclusive to human brains. And no human expert can doublecheck what the machine has done in a reasonable amount of time. Therefore there is nothing new about a machine performing a "cognitive" task that humans cannot match. Either the Singularity already happened in the 1950s or it is not clear what cognitive task would represent the coming of the Singularity.

To assess the progress in machine intelligence one has to show something (some intelligent task) that computers can do today that, given the same data, they could not have done fifty years ago. There has been a lot of progress in miniaturization and cost reduction, so that today it has become feasible to use computers for tasks for which we didn't use them fifty years ago; not because they were not intelligent enough to do them but because it would have been too expensive and it would have required several square kilometers of space. If that's "artificial intelligence", then we invented artificial intelligence when Konrad Zuse built the first computer in 1938. Today's computers can do a lot more things than the old ones just like new models of any machine (from kitchen appliances to mechanical reapers) can do a lot more things than old models. Incremental engineering steps lead to more and more advanced models for lower prices. Some day a company will

introduce coffee machines on wheels that can make the coffee and deliver the cup of coffee to your desk. And the next model will include voice recognition that understands "coffee please". Etc. This kind of progress has been going on since the invention of the first mechanical tool. It takes decades and sometimes centuries for the human race to fully take advantage of a new technology. The process of mastering that technology (of creating ever more sophisticated products based on that technology) is often what is referred to as "progress".

There is no question that progress has accelerated with the advent of electrical tools and further accelerated with the invention of computers. Whether these new classes of artifacts will eventually constitute a different kind of "intelligence" probably depends on your definition of "intelligence".

The way the Singularity would be achieved by intelligent machines is by these machines building more intelligent machines capable of building more intelligent machines and so forth. This would not be the first time that machines started building machines. The steam engine enabled mass production of steel, which in turn enabled the mass production of better steam engines, and this recursive loop continued for a while. James Watt himself, inventor of the steam engine that revolutionized the world, worked closely with John Wilkinson, who made the steel for Watt's engines using Watt's engines to make the steel. The human beings in this process can be viewed as mere intermediaries between machines that are evolving into better machines. The loop of machines building machines building (better) machines accelerated for a while. Eventually, the steam engine (no matter how sophisticated that accelerating positive feedback loop had made it) was made obsolete by a new kind of machine, the electrical motor. Again, electrical motors were used by manufacturers of motor parts that contributed to make better electrical motors used by manufacturers of electrical motor parts. We have been surrounded by machines that build better machines for a long time... but with human intermediaries designing and causing the quantum jump up.

Despite the fact that no machine has ever created another machine of its own will, and no software has ever created a software program of its own will, the Singularity crowd seems to have no trouble believing that a machine is coming soon created by a machine created by a machine and so forth, each generation of machines being smarter than the previous one.

i certainly share the concern that the complexity of a mostly automated world could get out of hand, something that has nothing to do with the degree of intelligence but just with the difficulty of managing complex systems. Complex, self-replicating systems that are difficult to manage have always existed. For example: cities, armies, post offices, subways, airports, sewers...

As far as i know, none of the people who make a living writing and talking about the coming Singularity have published a roadmap of predictions (when will machines be capable of just recognizing a written number with our accuracy? when will machines be able to cross a regular street without being run over by a car? when will machines be able to make improved versions of themselves?) that we can check in the coming years to see how credible these experts are.

A Look at the Evidence: A Comparative History of Accelerating Progress

A postulate at the basis of many contemporary books by futurists and self-congratulating technologists is that we live in an age of unprecedented rapid change and progress. Look closer and our age won't look so unique anymore.

As i wrote in the chapter titled "Regress" of my book "Synthesis", this perception that we live in an age of rapid progress is mostly based on the fact that we know the present much better than we know the past. One century ago, within a relatively short period of time, the world adopted the car, the airplane, the telephone, the radio and the record, while at the same time the visual arts went through Impressionism, Cubism and Expressionism, while at the

same time Quantum Mechanics and Relativity happened in science, while at the same time office machines (cash registers, adding machines, typewriters) and electrical appliances (dishwasher, refrigerator, air conditioning) dramatically changed the way people worked and lived, while at the same time Debussy, Schoenberg, Stravinsky and Varese changed the concept of music. The years since World War II have witnessed innovation that has been mostly gradual and incremental. We still drive cars (invented in 1886) and make phone calls (invented in 1876), we still fly on airplanes (invented in 1903) and use washing machines (invented in 1908), etc. Cars still have four wheels and planes still have two wings. We still listen to the radio and watch television. While the computer and Genetics have introduced powerful new concepts, and computers have certainly changed daily lives, i wonder if any of these "changes" compare with the notion of humans flying in the sky and of humans located in different cities talking to each other. There has been rapid and dramatic change before.

Does the revolution in computer science compare with the revolutions in electricity of a century ago? The smartphone and the Web have certainly changed the lives of millions of people, but didn't the light bulb, the phonograph, the radio and kitchen appliances change the world at least as much if not much more?

A history of private life in the last fifty years would be fairly disappointing: we wear pretty much the same clothes (notably T-shirts and blue jeans), listen to the same music (rock and soul were invented in the 1950s), run in the same shoes (sneakers date from the 1920s), and ride, drive and fly in the same kinds of vehicles (yes, even electric ones: Detroit Electric began manufacturing electric cars in 1907). Public transportation is still pretty much what it was a century ago: trams, buses, trains, subways. New types of transportation have been rare and have not spread widely: the monorail (that became reality with the Tokyo Monorail in 1964), the supersonic airplane (the Concorde debuted in 1976 but was retired in 2003), the magnetic levitation train (the Birmingham Maglev debuted in 1984, followed by Berlin's M-Bahn in 1991, but in

practice the Shanghai Maglev Train built in 2004 is the only real high-speed magnetic levitation line in service). The "bullet train" (widely available in Western Europe and the Far East since Japan's Shinkansen of 1964) is probably the only means of transportation that has significantly increased the speed at which people travel long distances in the last 50 years.

We chronically underestimate progress in previous centuries because most of us are ignorant about those eras. Historians, however, can point at the spectacular progress that took place in Europe during the Golden Century (the 13th century) when novelties such as spectacles, the hourglass, the cannon, the loom, the blast furnace, paper, the mechanical clock, the compass, the watermill, the trebuchet and the stirrup changed the lives of millions of people within a few generations; or the late 15th century when (among other things) the printing press enabled an explosive multiplication of books and when long-distance voyages to America and Asia created a whole new world.

Then one should discuss "change" versus "progress". Change for the sake of change is not necessarily "progress". Most changes in my software applications have negative, not positive effects, and we all know what it means when our bank announces "changes" in policies. If i randomly change all the cells in your body, i may boast of "very rapid and dramatic change" but not necessarily of "very rapid progress". Assuming that any change equates with progress is not only optimism: it is the recipe for ending up with exactly the opposite of progress. Out of the virtually infinite set of possible changes, only a tiny minority of them, a tiny subset, constitute progress.

The expression "exponential growth" is often used to describe our age, but trouble is that it has been used to describe just about every age since the invention of exponentials. In every age, there are always some things that grow exponentially, but others don't. For every technological innovation there was a moment when it spread "exponentially", whether it was church clocks or windmills, reading glasses or steam engines; and their "quality" improved

exponentially for a while, until the industry matured or a new technology took over. Moore's law is nothing special: similar exponential laws can be found for many of the old inventions. Think how quickly radio receivers spread: in the USA there were only five radio stations in 1921 but already 525 in 1923. Cars? The USA produced 11,200 in 1903, but already 1.5 million in 1916. By 1917 a whopping 40% of households had a telephone in the USA up from 5% in 1900. There were fewer than one million subscribers to cable television in 1984, but more than 50 million by 1989. The Wright brothers flew the first airplane in 1903: during World War I (1915-18) France built 67,987 airplanes, Britain 58,144, Germany 48,537, Italy 20,000 and the USA 15,000, for a grand total of almost 200 thousand airplanes; after just 15 years of its invention. Neil Armstrong stepped on the Moon in 1969, barely eight years after Yuri Gagarin had become the first human to leave the Earth's atmosphere.

Most of these fields then slowed down dramatically. And 44 years after the Moon landing we still haven't sent a human being to any planet and we haven't even returned to the Moon since the Apollo 17 in 1972. I am sure that similar statistics can be found for other old inventions, all the way back to the invention of writing. Perhaps each of those ages thought that growth in those fields would continue at the same pace forever. The wisest, though, must have foreseen that eventually growth starts declining in every field. Energy production increased thirteen-fold in the 20th century and freshwater consumption increased nine-fold, but today there are many more experts worried about a decline (relative to demand) than experts who believe in one more century of similar growth rates.

There has certainly been progress in telecommunications; but what difference does it make for ordinary people whether a message is sent in a split second or in two split seconds? In 1775 it took 40 days for the English public to learn that a revolution had started in the American colonies. Seven decades later it took minutes for the news of the Mexican War to travel to Washington.

That is real progress: from 40 days to a few minutes. The telegraph did represent "exponential" progress. Email, texting and chatting have revolutionized the way people communicate over long distances, but it is debatable whether that is (quantitatively and qualitatively) the same kind of revolution that the telegraph caused.

There are many "simpler" fields in which we never accomplished what we set out to accomplish originally, and pretty much abandoned the fight after the initial enthusiasm. We simply became used to the failure and forgot our initial enthusiasm. For example, domestic lighting progressed dramatically from gas lighting to Edison's light bulbs and Brush's arc lights of the 1880s and the first tungsten light-bulbs and then to the light-bulbs of the 1930s, but since then there has been very little progress: as everybody whose eyesight is aging knows too well, we still don't have artificial lighting that compares with natural sunlight, and so we need to wear reading glasses in the evening to read the same book that we can easily read during the day. A century of scientific and technological progress has not given us artificial lighting that matches sunlight.

I can name many examples of "change" that is often equated to "progress" when in fact it is not clear what kind of progress it is bringing. The number of sexual partners that a person has over a lifetime has greatly increased, and social networking software allows one to have thousands of friends all over the world, but i am not sure these changes (that qualify as progress from a strictly numerical point of view) result in happier lives. I am not sure that emails and text messages create the same bonds among people than the phone conversation, the letter on paper, the postcard and the neighbor's visit did.

One can actually argue that there is a lot of "regress", not "progress". We now listen to lo-fi music on computers and digital music players, as opposed to the expensive hi-fi stereos that were commonplace a generation ago. Mobile phone conversations are frequently of poor quality compared with the old land lines. We have access to all sorts of food 24 hours a day but the quality of that food is dubious. Not to mention "progress" in automated customer

support, which increasingly means "search for the answer by yourself on the Web" (especially from high-tech software giants like Microsoft and Google) as opposed to "call this number and an expert will assist you".

In the early days of the Internet (1980s) it was not easy to use the available tools but any piece of information on the Internet was written by very competent people. Basically, the Internet only contained reliable information written by experts. Today there might be a lot more data available, but the vast majority of what travels on the internet is: a) disinformation, b) advertising. It is not true that in the age of search engines it has become easier to search for information. Just the opposite: the huge amount of irrelevant and misleading data is making it more difficult to find the one webpage that has been written by the one great expert on the topic. In the old days her webpage was the only one that existed. (For a discussion on Wikipedia see the appendix).

Does the Internet itself represent true progress for human civilization if it causes the death of all the great magazines, newspapers, radio and television programs, the extinction of bookstores and record stores, and that it will be much rarer and harder to read and listen to the voices of the great intellectuals of the era? And all of this while at the same time massively increasing the power of corporations (via targeted advertising) and of governments (via systemic surveillance)? From the Pew Research Center's "State of the News Media 2013" report: "Estimates for newspaper newsroom cutbacks in 2012 put the industry down 30% since its peak in 2000. On CNN, the cable channel that has branded itself around deep reporting, produced story packages were cut nearly in half from 2007 to 2012. Across the three cable channels, coverage of live events during the day, which often require a crew and correspondent, fell 30% from 2007 to 2012... Time magazine is the only major print news weekly left standing".

Even the idea that complexity is increasing relies on a weak definition of "complexity". The complexity of using the many features of a smartphone is a luxury and cannot be compared with

the complexity of defending yourself from wild animals in the jungle or even with the complexity of dealing with weather, parasites and predators when growing food in a farm. The whole history of human civilization is a history of trying to reduce the complexity of the world. Civilization is about creating stable and simple lives in a stable and simple environment. By definition, what we call "progress" is a reduction in complexity, although to each generation it appears as an increase in complexity because of the new tools and the new rules that come with those tools. Overall, living has become simpler (not more complicated) than it was in the stone age. If you don't believe me, go and camp in the wilderness by yourself with no food and only stone tools.

In a sense, today's Singularity prophets assume that machine "intelligence" is the one field in which growth will never slow down, in fact it will keep accelerating.

Again, I would argue that it is not so much "intelligence" that has accelerated in machines (their intelligence is the same that Alan Turing gave them when he invented his "universal machine") but miniaturization. Moore's law (the one exponential law of electronic computing) has nothing to do with machine intelligence, but simply with how many transistors one can squeeze on a tiny integrated circuit. There is very little (in terms of intelligent tasks) that machines can do today that they could not have done in 1950 when Turing published his paper on machine intelligence. What has truly changed is that today we have extremely powerful computers squeezed into a palm-size smartphone at a fraction of the cost. That's miniaturization. Equating miniaturization to intelligence is like equating an improved wallet to wealth.

One has to decide which progress really matters for Artificial Intelligence: hardware or software? There has certainly been rapid progress in hardware technology (and in the science of materials in general) but the real question to me is whether there has been any real progress in software technology since the invention of binary logic and of programming languages. And a cunning software engineer would argue that even that question is not correct: there is

a difference between software engineering (that simply finds ways to implement algorithms in programming languages) and algorithms. The computer is a machine that executes algorithms. Anybody trying to create an intelligent machine using a computer is trying to find the algorithm or set of algorithms that will match or surpass human intelligence. Therefore it is neither progress in hardware not progress in software that really matters (those are simply enabling technologies) but progress in Computational Mathematics.

Kurzweil's book used a diagram titled "Exponential Growth in Computing" over a century, but i would argue that it is bogus because it starts with the electromechanical tabulators of a century ago: it is like comparing the power of a windmill with the power of a horse. Sure there is an exponential increase in power, but it doesn't mean that windmills will keep improving forever by the difference between horsepower and windpower. And it doesn't distinguish between progress in hardware and progress in software, nor between progress in software and progress in algorithms. What we would like to see is a diagram titled "Exponential Growth in Computational Math". As i am writing, most A.I. practitioners are looking for abstract algorithms that improve automatic learning techniques. Others believe that the correct way to achieve artificial intelligence should be to simulate the brain's structure and its neural processes, a strategy that greatly reduces the set of interesting algorithms. In that case, one would also want to see a diagram titled "Exponential Growth in Brain Simulation". Alas, any neurologist can tell you how far we are from understanding how the brain performs even the simplest daily tasks. Current brain simulation projects model only a small fraction of the structure of the brain, and provide only a simplified binary facsimile of it (neuronal states are represented as binary states, the variety of neurotransmitters is reduced to just one kind and there is no connection to a body), No laboratory has yet been able to duplicate the simplest brain we know, the brain of the 300-neuron roundworm: where's the exponential progress that would lead to a simulation of the 86

billion-neuron brain of Homo Sapiens (with its 100 trillion connections)? Since 1963 (when Sydney Brenner first proposed it), scientists worldwide have been trying to map the neural connections of the simplest roundworm, the *Caenorhabditis Elegans*, thus jump-starting a new discipline called Connectomics. So far they have been able to map only subsets of the worm's brain responsible for specific behaviors.

If you believe that an accurate simulation of brain processes will yield artificial intelligence (whatever your definition is of "artificial intelligence"), how accurate has that simulation to be? This is what neuroscientist Paul Nunez has called the "blueprint problem". Where does that simulation terminate? Does it terminate at the computational level, i.e. at simulating the exchanges of information within the brain? Does it terminate at the molecular level, i.e. simulating the neurotransmitters and the very flesh of the brain? Does it terminate at the electrochemical level, i.e. simulating electromagnetic equations and chemical reactions? Does it terminate at the quantum level, i.e. taking into consideration subatomic effects?

Kurzweil's "Law of Accelerating Returns" is nothing but the usual enthusiastic projection of the present into the future, a mistake made by millions of people all the time (millions of people buy homes when home values are going up believing that they would go up forever). Historically most technologies grew quickly for a while, then stabilized and continued to grow at a much slower pace until they became obsolete.

In general, the role of technology is sometimes overrated. Some increase in productivity is certainly due to technology, but in my opinion other contributions have been neglected too quickly. For example, Luis Bettencourt and Geoffrey West of the Santa Fe Institute have shown that doubling the population of a city causes on average an increase of 130% in its productivity ("A Unified Theory of Urban Living", 2010). This has nothing to do with technological progress but simply with urbanization. The rapid increase in productivity of the last 50 years may have more to do

with the rapid urbanization of the world than with Moore's law: in 1950 only 28.8% of the world's population lived in urban areas but in 2008 for the first time in history more than half of the world's population lived in cities (82% in North America, the most urbanized region in the world).

Predictions about future exponential trends have almost always been wrong. Remember the prediction that the world's population would "grow exponentially"? In 1960 Heinz von Foerster predicted that population growth would become infinite by Friday the 13th of November 2026. Now we are beginning to fear that it will actually start shrinking (it already is in Japan and Italy). Or the prediction that energy consumption in the West will grow exponentially? It peaked a decade ago; and, as a percentage of GDP, it is actually declining rapidly. Life expectancy? It rose rapidly in the West between 1900 and 1980 but since then it has barely moved. War casualties were supposed to grow exponentially with the invention of nuclear weapons: since the invention of nuclear weapons the world has experienced the lowest number of casualties ever, and places like Western Europe that had been at war nonstop for 1500 year have not had a major war in 60 years.

There is one field in which i have witnessed rapid (if not exponential) progress: Genetics. This discipline has come a long way in just 70 years, since Oswald Avery and others identified DNA as the genetic material (1944) and James Watson and Francis Crick discovered the double-helix structure of DNA (1953). Frederick Sanger produced the first full genome of a living being in 1977, Kary Banks Mullis developed the polymerase chain reaction in 1983, Applied Biosystems introduced the first fully automated sequencing machine in 1987, William French Anderson performed the first procedure of gene therapy in 1990, Ian Wilmut cloned a sheep in 1997, the sequencing of the human genome was achieved by 2003, and Craig Venter and Hamilton Smith reprogrammed a bacterium's DNA in 2010. The reason that there has been such dramatic progress in this field is that a genuine breakthrough happened with the discovery of the structure of DNA. I don't believe

that there has been an equivalent discovery in the field of Artificial Intelligence.

Economists would love to hear that progress is accelerating because it has an impact on productivity, which is one of the two factors driving GDP growth. GDP growth is basically due to population growth plus productivity increase. Population growth is coming to a standstill in all developing countries (and declining even in countries like Iran and Bangladesh) and, anyway, the biggest contributor to workforce growth was actually women, which came to the workplace by the millions, but now that number has stabilized. If progress were accelerating, you'd expect productivity growth to accelerate. Instead, despite all the hoopla about computers and the Internet, productivity growth of the last 30 years has averaged 1.3% compared with 1.8% in the previous 40 years. Economists like Jeremy Grantham now predict a future of zero growth. Not just deceleration but a shrieking halt.

Whenever i meet someone who strongly believes that machine intelligence is accelerating under our nose, i ask him/her a simple question: "What can machines do today that they could not do five years ago?" If their skills are "accelerating" and within 20-30 years they will have surpassed human intelligence, it shouldn't be difficult to answer that question. So far the answers to that question have consistently been about incremental refinements (e.g., the new release of a popular smartphone that can take pictures at higher resolution) and/or factually false ("they can recognize cats", which is not true because i have seen no such machine, nor heard of any announcement of any such machine coming to the market any time soon).

In 1939 at the World's Fair in New York the General Motors Futurama exhibit showed how life would be in 1960 thanks to technological progress: the landscape was full of driverless cars. The voiceover said: "Does it seem strange? Unbelievable? Remember, this is the world of 1960!" Twentyone years later the world of 1960 turned out to be much more similar to the world of 1939 than to the futuristic world of that exhibit.

On the 3rd of April 1988 the Los Angeles Times Magazine ran a piece titled "L.A. 2013" in which experts predicted how life would look like in 2013 (the year when I am writing this piece). They were comfortable predicting that the average middle-class family would have two robots to carry out all household chores including cooking and washing, that kitchen appliances would be capable of intelligent tasks, and that people would commute to work in self-driving cars. How many robots do you have in your home and how often do you travel in a self-driving car in 2013?

In 1964 Isaac Asimov wrote an article in the New York Times (August 16) titled "Visit to the World's Fair of 2014" in which he predicted what the Earth would look like in 2014. He envisioned that by 2014 there would be Moon colonies and all appliances would be cordless.

The future is mostly disappointing. As Benjamin Bratton wrote in December 2013: "So little of the future promised in TED talks actually happens".

People who think that progress has been dramatic are just not aware of how fast progress was happening before they were born and of how high the expectations were and of how badly those expectations have been missed by current technology.

Intermezzo: In Defense of Regress

We accept as "progress" many innovations whose usefulness is dubious at best.

Any computer with a "mouse" requires the user to basically have three hands.

Watching films on digital media such as DVDs is more difficult for a professional critic than watching them on videotapes because stopping, rewinding, forwarding and, in general, pinpointing a scene is much easier and faster on analog tapes than on digital files.

Computer's optical disc drives and car's radios that you have to push (instead of pull) in order to open are simply more likely to break and don't really add any useful feature.

The automation of cash registers means that it takes longer to pay than to find the item you want to buy (and you cannot buy it at all if the cash register doesn't work).

The car keys with an embedded microchip cost 140 times more to duplicate than the old chip-less car keys (as i just found out at the local Toyota dealer).

Most portable gadgets used to operate with the same AA or AAA batteries. When on the road, you only had to worry about having those spare batteries. Now most cameras work only with proprietary rechargeable batteries: the fact that they are "rechargeable" is useless if they die in a place where you cannot recharge them, which is the case whenever you are far from a town or forgot the charger at home. I don't see this as progress compared with the cheap, easily replaceable AA batteries that i could also use with my hiking GPS, my headlight and my walkie-talkie. In fact, Nikon mentions it as a plus that its Coolpix series is still "powered by readily available AA batteries".

It is hard to believe that there was a time when you would pick up the phone and ask an operator to connect you to someone. Now you have to dial a 10-digit number, and sometimes a 13-digit number if you are calling abroad. More recently there used to be telephone directories to find the phone number of other telephone subscribers. I remember making fun of Moscow when we visited it in the 1980s because it didn't have a telephone directory. In the age of mobile phones the telephone directory has disappeared: you can know a subscriber's number only if someone gives it to you. Apparently the Soviet Union was the future, not the past.

Never since the 1950s have phone communications been so rudimentary as after the introduction of the mobile phone. Conversations invariably contain a lot of "Can you hear me?" like in the age of black and white movies.

I felt relieved in a Mexican town where there was a public phone at every corner: drop a coin and you are making a phone call. Wow. No contract needed, and no "can you hear me?"

Mobile phone ringers that go off in public places such as movie theaters and auditoria (and that obnoxiously repeat the same music-box refrain in a mechanical tone) do not improve the experience.

Voice recognition may have represented an improvement when it allowed people to say numbers rather than press them on the phone keyboard; but now the automated system on the other side of the phone asks you for names of cities or even “your mother’s maiden name”, and never gets them right (especially if you, like me, have an accent), or for long numbers (such as the 16-digit number of my credit card) that you have to repeat over and over again.

Thanks to air-conditioned buildings with windows that are tightly sealed, we freeze in the summer and sometimes catch bronchitis.

When you travel by bus or train, you need to get your ticket at a machine or have exact change to buy it on the bus, hardly an improvement over the old system of paying the conductor when you board. New buses and trains are climatized: it is impossible to take decent pictures of the landscape because the windows cannot be opened and are dimmed.

Talking of windows, the electric windows of your car won’t operate if the car’s battery dies (the old “roll down the window” does not apply to a car with dead battery).

Printing photographs has become more, not less, expensive with the advent of digital cameras, and the quality of the print is debatable.

The taximeter, rarely used in developing countries but mandatory in “advanced” countries, has always been a favorite of mine. Basically, a taxi driver asks you to buy a good without telling you the price until you have already used the good and you cannot change your mind. The taximeter often increases the cost of a ride because you can’t bargain anymore as you would normally do based on the law of supply and demand (for example, in situations when the taxi driver has no hopes for other customers). Furthermore, the taximeter motivates unscrupulous taxi drivers to take the longest and slowest route to your destination, whereas a

negotiated price would motivate the driver to deliver you as quickly as possible.

Websites with graphics, animation, pop-up windows, “click here and there” cause you to spend most of the time “navigating” instead of reading the information you were looking for.

We the consumers passively accept too many of these dubious “improvements”.

Most of these “improvements” may represent progress, but the question is “progress for whom”? Pickpockets used to steal one wallet at a time. The fact that today a hacker can steal millions of credit card numbers in an instant constitutes progress, but progress for whom?

We live in a world of not particularly intelligent machines, but certainly in a world of machines who like to beep. My car beeps when I start it, when I leave the door open, and, of course, if I don't fasten my seat belt. Note that it doesn't beep if something much more serious happen, like the alternator dies or the oil level is dangerously low. My microwave oven beeps when it's done, and it keeps beeping virtually forever unless someone opens its door (it doesn't matter that you actually pick up the food, just open the door). My printer beeps when it starts, when I feed it paper, and when something goes wrong (a blinking message on the display is not enough, apparently). Best of all, my phone beeps when I turn it off, and, of course, sometimes I turn it off because I want it silent: it will beep to tell everybody that is being turned silent. I think that every manual should come with instructions on how to disable the beeping device: “First and foremost, here is how you can completely shut up your device once and forever”.

Last but not least, something is being lost in the digital age, something that was the fundamental experience of (broadly defined) entertainment. During a vacation in a developing country i watched as a girl came out of the photographer's shop. She couldn't wait and immediately opened the envelope that contained her photographs. I witnessed her joy as she flipped through the pictures. The magic of that moment, when she sees how the pictures came, will be gone

the day she buys her first digital camera. There will be nothing special about watching the pictures on her computer's screen. There will be no anxious waiting while she uploads them to the computer because, most likely, she will already know how the pictures look like before she uploads them. Part of the magic of taking photographs is gone forever, replaced by a new, cold experience that consists in refining the photograph with digital tools until it is what you want to see, and then posting it on social media in an act of vanity.

Or take live events. The magic of a live sport event used to be the anxious wait for the event to start and then the "rooting" for one of the competitors or teams. In the age of TiVo, one can watch a "live" event at any time by conveniently "taping" it. Many live events are actually broadcasted with a slight delay (you may find on the Internet the result of a soccer game that is still going on according to your television channel). Thus the "waiting" and the "rooting" are no longer the two fundamental components of the "live" experience. The whole point of watching a live event was the feeling you're your emotional state could somehow influence the result. If the event is recorded (i.e., already happened), that feeling disappears and you have to face the crude reality of your impotence to affect the result. But then what's the point of rooting? Thus the viewer is unlikely to feel the same emotional attachment to the game that s/he is watching. In the back of her/his mind, it is clear that the game has already finished. The experience of watching the "live" event is no longer one of anxiety but one of appreciation. Told by a friend that it was a lousy game, the viewer may well decide not to watch the taped event. These are examples of how the digital appliances are altering not only the syntax of daily life, but its very semantics.

Yes, i know that Skype and Uber and many new services can solve or will solve these problems, but the point is that these gadgets and features were conceived and understood as "progress" when they were introduced (and usually amid much fanfare). The very fact that platforms such as Skype and Uber have been successful proves that the quality of services in those fields had

overall regressed, not progressed, and therefore there was an opportunity for someone to restore service to a decent level.

We should always pause and analyze whether something presented as "progress" truly represents progress. And for whom.

The Feasibility of Superhuman Intelligence

Maybe the Singularity argument is mostly a philosophical (not scientific) argument.

What are the things that a superhuman intelligence can do and i cannot do? If the answer is "we cannot even conceive them", then we are back to the belief that angels exist and miracles happen, something that eventually gave rise to organized religions.

The counter-argument is that mono-cellular organisms could not foresee the coming of multicellular organisms. Maybe. But bacteria are still around, and probably more numerous than any other form of life in our part of the universe. The forms of life that came after bacteria were perhaps inconceivable by bacteria but, precisely because they were on a different plane, they hardly interact. We kill bacteria when they harm us but we also rely on many of them to work for us (our body has more bacterial cells than human cells). In fact, one could claim that a superhuman intelligence already exists, and it's the planet as a whole, Gaia, of which we are just one of the many components.

On one hand, superhuman intelligence should exist because of the "cognitive closure", a concept popularized by Colin McGinn in "The Problem Of Consciousness" (1991). The general idea is that every cognitive system (e.g., every living being) has a "cognitive closure": a limit to what it can know. A fly or a snake cannot see the world the way we see it because they do not have the same visual system that humans have. In turn, we can never know how it feels to be a fly or a snake. A blind person can never know what "red" is even after studying everything that there is to be studied about it.

According to this idea, each brain (including the human brain) has a limit to what it can possibly think, understand, know. In particular, the human brain has a limit that will preclude humans from understanding some of the ultimate truths of the universe. These may include spacetime, the meaning of life, and consciousness itself. There is a limit to how "intelligent" we humans can be. According to this view, there should exist cognitive systems that are "superhuman" in that they don't have the limitations that our cognition has.

On the other hand, i am not sure if we (humans) can intentionally build a cognitive system whose cognitive closure is larger than ours, i.e., a cognitive system that can "think" concepts that we cannot think. It sounds a bit of a contradiction in terms that a lower form of intelligence can intentionally build a highest form of intelligence. However, it is not a contradiction that a lower form of intelligence can accidentally (by sheer luck) create a higher form of intelligence.

That is an argument in favor of the feasibility of superhuman intelligence. Alas, there are also arguments against such feasibility. Notably, David Deutsch in "The Beginning of Infinity" (2011) argues that there is nothing in our universe that the human mind cannot understand, as long as the universe is driven by universal laws. I tend to agree with Colin McGinn that there is a "cognitive closure" for any kind of brain, that any kind of brain can only do certain things, and that our cognitive closure will keep us from ever understanding some things about the world (perhaps the nature of consciousness is one of them); but in general i agree with Deutsch: if something can be expressed in formulas, then we will eventually "discover" it and "understand" it; and, if everything in nature can be expressed in formulas, then we (human intelligence) will eventually "understand" everything, i.e. we are the highest form of intelligence that can possibly exist. So the only superhuman machine that would be too intelligent for humans to understand is a machine that does not obey the laws of nature, i.e. is not a machine.

Intelligence is the capability to understand the universe, everything that exists in the universe. I am not sure what

"superhuman" intelligence would be. (Counterargument: of course you don't know, because you are only human).

The definition of intelligence is vague, and it is vague because such a thing does not exist. We tend to call "intelligence" the whole repertory of human skills, from eating to theorizing. It would be better to break down human life into skills and then, for each skill, to assess how far we are from having machines that perform those skills. If, instead, "artificial intelligence" refers to only a subset of those skills, it would be nice to list them. I suspect that 100 researchers in Artificial Intelligence would come up with 100 different lists, and that 100 researchers in Singularity matter would come up with the most vague of lists.

The Singularity crowd also brushes aside as irrelevant the old arguments against Artificial Intelligence. The debate has been going on for 50 years, and it has never been settled. In 1935 Alonzo Church proved a theorem, basically an extension of Kurt Goedel's incompleteness theorem (1931) to computation, that first order logic is undecidable. Similarly, in 1936 Alan Turing proved that the "halting problem" is undecidable for Universal Turing Machines (it cannot be proven whether a computer will always find a solution to every problem). These are not details because these are theorems that humans can understand, and in fact humans discovered them and proved them, but machines cannot. Critics of Artificial Intelligence argue that machines will never become as intelligent as humans for the same reason that $1+1$ is not 3 but 2.

Even brushing these logical arguments aside (which is like brushing aside the logical evidence that the Earth revolves around the Sun), one has to wonder for how long it will make sense to ask the question whether superhuman intelligence is possible. If the timeframe for intelligent machines is centuries and not decades like the optimists believe, then it's like asking an astronaut "Will it at some point be possible to send a manned spaceship to Pluto?" Yes, it may be very possible, but it may never happen: not because it's impossible but simply because we may invent teleportation that will make spaceships irrelevant. Before we invent intelligent

machines, synthetic biology or some other discipline might have invented something that will make robots and the likes irrelevant. The timeframe is not a detail either.

What would make me believe in the coming of some truly intelligent machine is a conceptual breakthrough in Artificial Intelligence. That might come as neuroscientists learn more about the brain; or perhaps a Turing-like mathematician will come up with some new way to build intelligent machines. But today we still use Turing Machines and binary logic, exactly like in 1946.

There is another line of attack against superhuman intelligence. The whole idea of the Singularity is based on a combination of the cognitive closure and the old fallacy of assuming that human progress (the progress of the human mind) ended with you. Even great philosophers like Georg Hegel fell for that fallacy. If we haven't reached the cognitive closure yet, then there is no reason why human intelligence should stall. Today (2013) that perception is partially justified because progress has slowed down compared with a century ago: anybody who has studied the history of Quantum Mechanics and Relativity has been amazed by the incredible (literally incredible) insight that those scientists had. The fallacy consists in believing that the human mind has reached a maximum of creativity and will never go any further. We build machines based on today's knowledge and creativity skills. Those machines will be able to do what we do today except for one thing: the next step of creativity that will make us think in different ways and invent things (not necessarily machines) of a different kind. For example, we (humans) are using machines to study synthetic biology, and progress in synthetic biology may make many machines disposable. Today's electronic machines may continue to exist and evolve, just like windmills existed and evolved and did a much better job than humans at what they were doing, and machines might even build other machines, but in the future they might be considered as intelligent as the windmills. Potentially, there is still a long way to go for human creativity. The Singularity crowd cannot imagine the future of human intelligence the same way that

someone in 1904 could not imagine Relativity and Quantum Mechanics.

What is the Opposite of the Singularity?

What worries me most is not the rapid increase in machine intelligence but the decrease in human intelligence.

The Turing Test cannot be abstracted from a sociological context. Whenever one separates sociology and technology, one misses the point. The Turing Test is commonly understood as: when can we say that a machine has become as intelligent as humans? But the Turing Test is about humans as much as it is about machines because it can equivalently be formulated as: when can we say that humans have become less intelligent than a machine? The implicit assumption was that achieving the Turing Point (the point when the machine has become as smart as humans) requires a smarter machine. But there's another way to pass the Turing Test: make dumber humans.

Humans have always become dependent on the tools they invented. When they invented writing, they lost memory skills. On the other hand, they gained a way to store a lot more knowledge and to broadcast it a lot faster. Ditto for all other inventions in history: a skill was lost, a skill was acquired. We cannot replay history backwards and we will never know what the world would be like if humans had not lost those memory skills. Indirectly we assume that the world as it is now is the best that it could have been. In reality, over the centuries the weaker memory skills have been driving an explosion of tools to deal with weak memory. Each tool, in turn, caused the decline of another skill. It is debatable if the invention of writing was worth this long chain of lost skills. This process of "dumbization" has been going on throughout society and accelerated dramatically ("exploded") with the digital devices. The computer caused the decline of calligraphy. Voice recognition will cause the decline of writing. In a sense, technology is about giving

dumb people the tools to become dumber and still continue to live a happy life.

What can machines do now that they could not do 50 years ago? They are just faster, cheaper and can store larger amounts of information. These factors made them ubiquitous. What could humans do 50 years ago that they cannot do now? Ask your grandparents and the list is very long, from multiplication to orientation, from driving in chaotic traffic to fixing a broken shoe. Or just travel to an underdeveloped country where people still live like your old folks used to live and you will find out how incapable you are of simple actions that are routine for them. When will we see a robot that is capable of crossing a street with no help from the traffic light? It will probably take several decades. When will we get to the point that the average person is no longer capable of crossing a street without help from the traffic light? That day is coming much sooner. Judging from simple daily chores, human intelligence appears to be not "exploding" but imploding. Based on the evidence, one can argue that machines are not getting much smarter (just faster), while humans are getting dumber; hence very soon we will have machines that are smarter than humans but not only because machines got smarter.

The age of digital devices is enabling the average person to have all sorts of knowledge at her fingertips. That knowledge originally came from someone who was "intelligent" in whichever field. Now it can be used by just about anybody who is not "intelligent" in that field. This user has no motivation to actually learn it: she can just "use" somebody else's "intelligence". The "intelligence" of the user decreases, not increases (except, of course, for the intelligence on how to operate the devices, but, as devices become easier and easier to use, eventually the only intelligence required will be to press a button to turn the device on). Inevitably, humans are becoming ever more dependent on machines, and less viceversa.

I chair/organize/moderate cultural events in the Bay Area and, having been around in the old days of the overhead projectors, I'm incredulous when a speaker cannot give her/his talk because

her/his computer does not connect properly to the room's audio/visual equipment and therefore s/he cannot use the prepared slide presentation. For thousands of years humans were perfectly capable of giving a talk without any help from technology. Not anymore, apparently.

The Turing Test could be a self-fulfilling prophecy: at the same time that we (claim to) build "smarter" machines, we are making dumber people.

My concern is not for machines that are becoming too intelligent, but for humans who are becoming less intelligent. What might be accelerating is the loss of human skills. Every tool deprives humans of the training they need to maintain a skill (whether arithmetic or orientation) and every interaction with machines requires humans to lower their intelligence to the intelligence of machines (e.g., to press digits on a phone in order to request a service). We can argue forever if the onboard computer of a self-driving car is really "driving", but we know for sure what the effect of self-driving cars will be: to raise a generation of humans that are not capable of driving anymore. Every machine that replaces a human skill (whether the pocket calculator or the street navigator) reduces the training that humans get in performing that skill (such as arithmetic and orientation), and therefore causes humans to lose that skill. This is an ongoing experiment on the human race that is likely to have a spectacular result: the first major regression in intelligence in the history of our species.

To be fair, it is not technology per se that makes us dumber. The same system that produces technology also makes us dumber. The first step usually consists in some rules and regulations that simplify and normalize a process, whether serving food at a fast-food chain or inquiring about the balance of your checking account or driving a car. Once those rules and regulations are in place, it gets much easier to replace human skills with technology: the human skills required to perform those tasks have been reduced dramatically, and, in that sense, humans have become "dumber" at those tasks. In a sense, technology is often an effect, not a cause: once the

skills required to perform a task have been greatly downgraded, it is quite natural to replace the human operator with a machine.

Paraphrasing something Bertrand Russell said about Ludwig Wittgenstein, we are weary of thinking and we are building a society that would make such an activity unnecessary. Then, of course, an unthinking machine would equal an unthinking human, not because the machine has become as thinking as the human, but because the human has become as unthinking as the machine.

The Turing Test can be achieved in two ways: 1. by making machines so intelligent that they will seem human; 2. by making humans so stupid that they will seem mechanical.

With all due respect, when i interact with government officials or corporate employees, the idea that these people, trained like monkeys to repeatedly say and do the prescribed routine, represent a species that is heading towards a higher form of intelligence seems simply laughable.

What will "singular" mean in a post-literate and post-arithmetic world?

Intermezzo: The Attention Span

This has more to do with modern life than with machines, but it is related to the idea of an "intelligence implosion".

I worry that the chronic scarcity of time of our age is pushing too many decision makers to take decisions having heard only very superficial arguments. The "elevator pitch" has become common even in academia. A meeting that lasts more than 30 minutes is a rarity (in fact, a luxury from the point of view of the most powerful, and therefore busiest, executives). You can't get anybody's attention for more than 20 minutes, but some issues cannot be fully understood in 20 minutes; and some great scientists are not as good at rhetorical speech as they are at their science, which means that they may lose a 20-minute argument even if they are 100% right. Too many discussions are downgraded because they take

place by texting on smartphones, whose tiny keyboards discourage elaborate messages. Twitter's 140-character posts are emblematic of the shrinking attention span. I am not afraid of the human race losing control of its machines as much as i am afraid that the human race will self-destruct because of the limitations of the "elevator pitch" and of the "tweet"; because of the chronic inability of decision makers, as well as of the general public, to fully understand an issue. (Incidentally, the reason that we have fewer and fewer investigative reporters in news organizations is the same, i.e. the reduced attention span of the readers/viewers, with the result that the reliability of news media is constantly declining). It has become impossible to properly organize a hiking trip because the participants, accustomed to tweets and texting, will only read the first few lines of a lengthy email; but, if you didn't prepare properly, you might get into serious trouble on a big mountain, perhaps even die. Now multiply this concept a few billion times to get the dimension of humanity's major problems, and you should understand why the last of my concerns is that machines may become too intelligent and the first of my concerns is that humans might become too dumb. As i type this sentence, Elon Musk (at MIT's AeroAstro 100 conference in October 2014) and others are worried that machines may get so smart that they will start building smarter machines; instead, i am worried that people's attention span is becoming so short that it will soon be impossible to explain the consequences of a short attention span. I don't see an acceleration in what machines can do, but i do see a deceleration in human attention... if not in human intelligence in general).

To clarify, there are three ways that we can produce "dumber" humans. All three are related to technology but in opposite ways. First, there is the simple fact that a new technology makes some skills irrelevant, and within one generation those skills are lost. Pessimists argue that little by little we become less human. Optimists claim that the same technology enables new skills to develop. I can personally attest that both camps are right: the computer and email have turned me into a highly-productive multi-

tasking cyborg, and at the same time they have greatly reduced my skills in writing polite and touching letters to friends and relatives (with, alas, parallel effects on the quality of my poetry). The pessimists think that the gains do not offset the losses (the "dumbization"), especially when it comes to losing basic survival skills. Secondly, the rules and regulations that society introduces for the purpose of making us safer and more efficient end up making us think less and less, i.e. behave more and more like (non-intelligent) machines. Thirdly, the frantic lives of overworked individuals have greatly reduced their attention span, which may result in a chronic inability to engage in serious discussions, i.e. in a more and more superficial concept of "intelligence", i.e. in the same limited cognitive experience of lower forms of life.

Semantics

In private conversations about "machine intelligence" i like to quip that it is not intelligent to talk about intelligent machines: whatever they do is not what we do, and, therefore, is neither "intelligent" nor "stupid" (attributes invented to define human behavior). Talking about the intelligence of a machine is like talking about the leaves of a person: trees have leaves, people don't. "Intelligence" and "stupidity" are not properties of machines: they are properties of humans.

Machines don't think, they do something else. Machine intelligence is as much an oxymoron as human furniture. Machines have a life of their own, but that "life" is not human life.

We apply to machines many words invented for humans simply because we don't have a vocabulary for the states of machines. For example, we buy "memory" for our computer, but that is not memory at all: it doesn't remember (it simply stores) and it doesn't even forget, the two defining properties of (biological) memory. We call it "memory" for lack of a better word. We talk about the "speed"

of a processor but it is not the "speed" at which a human being runs or drives. We don't have the vocabulary for machine behavior. We borrow words from the vocabulary of human behavior. It is a mistake to assume that, because we use the same word to name them, then they are the same thing. If i see a new kind of fruit and call it "cherry" because there is no word in my language for it, it doesn't mean it is a cherry. A computer does not "learn": what it does when it refines its data representation is something else (that we don't do).

It is not just semantics. Data storage is not memory. Announcements of exponentially increasing data storage miss the point: that statistical fact is as relevant to intelligence as the exponential increase in credit card debt. Just because a certain sequence of zeroes and ones happens to match a sequence of zeroes and ones from the past it does not mean that the machine "remembered" something. Remembering implies a lot more than simply finding a match in a data storage. Memory does not store data. In fact, you cannot retell a story accurately (without missing and possibly distorting tons of details) and you cannot retell it twice with the same words (each time you will use slightly different words). Ask someone what her job is, something that she has been asked a thousand times, and she'll answer the question every time with a different sequence of words, even if she tries to use the same words she used five minutes earlier. Memory is "reconstructive". We memorize events in a very convoluted manner, and we retrieve them in an equally convoluted manner. We don't just "remember" one thing: we remember our entire life whenever we remember something. It's all tangled together. You understand something not when you repeat it word by word like a parrot (parrots can do that, and tape recorders can do that) but when you summarize it in your own words, different words than the ones you read or heard: that is what we call "intelligence". I am always fascinated, when i write something, to read how readers rewrite it in their own words, sometimes using completely different words, and sometimes saying it better than i did.

It is incredible how bad our memory is. A friend suggested an article by David Carr about Silicon Valley published in the New York Times "a few weeks ago". It took several email interactions to figure out that a) the article was written by George Packer, b) it was published by the New Yorker, c) it came out one year earlier. And, still, it is amazing how good our memory is: it took only a few sentences during a casual conversation for my friend to relate my views on the culture of Silicon Valley to an article that she had read one year earlier. Her memory has more than just a summary of that article: it has a virtually infinite number of attributes linked to that article such that she can find relevant commonalities with the handful of sentences she heard from me. It took her a split second to make the connection between some sentences of mine (presumably ungrammatical and incoherent sentences because we were in a coffee house and I wasn't really trying to compose a speech) and one of the thousands of articles that she has read in her life.

All forms of intelligence that we have found so far use memory, not data storage. I suspect that, in order to build an artificial intelligence that can compete with the simplest living organism, we will first need to create artificial memory (not data storage). Data storage alone will never get you there, no matter how many terabytes it will pack in a millimeter.

What computers do is called "savant syndrome" in the scientific literature: idiots (very low intelligence quotient) with a prodigious memory.

Data is not knowledge either: having amassed all the data about the human genome does not mean that we know how human genes work. We know a tiny fraction of what they do even though we have the complete data.

I was asking a friend how the self-driving car works in heavy traffic and he said "the car knows which other cars are around". I object that the car does not "know" it. There is a system of sensors that continuously relay information to a computer that, in turn, calculates the trajectory and feeds it into the motor controlling the steering

wheel. This is not what we mean when we say that we "know" something. The car does not "know" that there are other cars around, and it does not "know" that cars exist, and it doesn't even "know" that it is a car. It is certainly doing something, but what it is doing is something else than "knowing". This is not just semantics: because the car does not "know" that it is a car surrounded by other cars driving on a road, it also lacks all the common sense or general knowledge that comes with that knowledge. If an elephant fell from the sky, a human driver would be at least surprised (and probably worried of stranger things ahead), whereas a car would simply interpret it as a car parked in the middle of the highway. When a few months ago (2013) Stanford researchers trained a robot to take the elevator, they realized that there was a non-trivial problem: the robot stops in front of the glass doors of the elevator interpreting its own reflection into it as another robot. The robot does not "know" that the thing is a glass door otherwise it would easily realize that there is no approaching robot, just a reflection getting bigger like all reflections do when you walk towards a mirror.

It is easy to claim that, thanks to Moore's law, today's computers are one million times faster than the computers of the 1980s, and that a smartphone is thousands of times faster than the fastest computer of the 1960s. But faster at what? Your smartphone is still slower than a snail at walking (it doesn't move, does it?) and slower than television at streaming videos. Plenty of million-year old artifacts are faster than the fastest computer at all sorts of biological processes. And plenty of analog devices (like television) are still faster than digital devices at what they do. If Moore's law applies to the next 20 years, there will be another million times improvement in processing speed and storage capacity. But that may not increase at all the speed at which a computer can summarize a film. The movie player that i use today on my laptop is slower (and a lot less accurate) in rewinding a few scenes of the film than the old videotape player of twenty years ago, no matter how "fast" the processor of my laptop is.

We tend to use cognitive terms only for machines that include a computer, and that started way back when computers were invented. Then the cognitive vocabulary tempts people to attribute "states of mind" to those machines. We don't usually do this to other machines. A washing machine washes clothes. If a washing machine is introduced that washes tons of clothes in a split second, consumers will be ecstatic, but presumably nobody would talk about it as if it related to a human or superhuman intelligence. And note that appliances do some pretty amazing things. There's even a machine called "television set" that shows you what is happening somewhere else, a feat that no intelligent being can do. We don't attribute cognitive states to a television set even though the tv set can do something that requires more than human intelligence.

Take happiness instead of intelligence, since "intelligence" is now so often related to machines. One of the fundamental states of human beings is "happiness". When is a machine "happy"? The question is meaningless: it's like asking when does a human being need to be watered? You water plants, not humans. Happiness is a meaningless word for machines. Some day we may start using the word "happy" to mean, for example, that the machine has achieved its goal or that it has enough electricity; but it would simply be a linguistic expedient. The fact that we may call it "happiness" does not mean that it "is" happiness. If you call me Peter because you can't spell my name, it does not mean that my name is Peter.

Note that the most intelligent machines ever built, such as IBM's "Deep Blue" and "Watson", are totally useless: they are not used for anything other than the specific task that made them famous. They can't even cook an omelette, they cannot drive a car, they cannot sit on the sidewalk and gossip about the neighborhood (i am thinking of human activities that we normally don't consider "very intelligent"). A very dumb human being can do a lot more than the smartest machines ever built, and that's probably because there is a fundamental misunderstanding about what "intelligent" means.

Non-human Intelligence is Already Here

There are already many kinds of intelligence that we cannot match nor truly comprehend. Bats can avoid objects in absolute darkness at impressive speeds and even capture flying insects because their brain is equipped with a high-frequency sonar system. Migratory animals can orient themselves and navigate vast territories without any help from maps. Birds are equipped with a sixth sense for the Earth's magnetic field. Some animals have the ability to camouflage. The best color vision is in birds, fish, and some insects. Many animals have night vision. Animals can see, sniff and hear things that we cannot, and airports still routinely employ sniffing dogs (not sniffing humans) to detect food, drugs and explosives. And don't underestimate the brain of an insect either: how many people can fly and land upside down on a ceiling?

Virtually all dogs existing today are artificial living beings: they are the result of selective breeding strategies. If you think that your dog is intelligent, then you have "artificial intelligence" right at home.

Ironically, when Deborah Gordon discovered that ant colonies use a packet-switching technique very similar to the one employed by the Internet ("The Regulation of Ant Colony Foraging Activity without Spatial Information", 2012), the media wrote that ants can do what the Internet does when in fact ants have been doing it for about 100 million years: it took human intelligence 100 million years to figure out the same system of communication devised by ant intelligence.

Summarizing, many animals have powers we don't have. We have arbitrarily decided that any skill possessed by other animals and not by humans is an inferior skill, whereas any skill possessed by humans and not by other animals is a superior skill. This leads me to wonder what will make a skill "superhuman": just the fact that it is possessed by a machine instead of an animal?

And, of course, we already built machines that can do things that are impossible for humans. The clock, invented almost a thousand years ago, does something that no human can do: keeping time.

Telescopes and microscopes can see things that humans cannot see. We can only see a human-level rendition by those machines, which is equivalent to a higher intelligence explaining something in simpler terms to a lower intelligence. We cannot do what light bulbs do. We cannot touch the groove of a rotating vinyl record and produce the sound of an entire philharmonic orchestra. And, of course, one such appliance is the computer, that can perform calculations much faster than any mathematician could. Even the pre-digital calculators of the 1940s (for example, the ones used to calculate ballistic trajectories) could calculate faster than human brains. In fact, we have always been post-human, coexisting with, and relying on, and being guided by, technology that was capable of super-human feats (and there have always been philosophers debating whether that post-human condition is anti-human or pro-human).

The intelligence of both animals and tools is not called "superhuman" simply because we are used to it. We are not used to robots doing whatever it is that they will do better than us and therefore we call it "superhuman" when in fact we should just call all of these "non-human life"; and maybe "non-human intelligence" depending on your definition of "intelligence".

If a machine ever arises (and proliferates) that is alive but does things that we cannot do, it will just be yet another form of non-human life: not the first one, not the last one. Of course, there are plenty of forms of life that are dangerous to humans, mostly very tiny ones (like viruses and ticks). It comes with the territory. If you want to call it "superhuman", suit yourself.

One gene can make a huge difference in brain structure and function, as the tiny difference between the chimp's DNA and human DNA proves. Gene therapy is already here and that is indeed progressing quickly. Changing the genes of the human DNA may have consequences that are orders of magnitudes bigger than we can imagine. That is one of the reasons why I tend to believe that "superhuman" intelligence is more likely to come from synthetic biology than from computers.

There are even qualitative differences in the "intelligences" of a person as the person grows and changes. Psychologists since at least Jean Piaget have studied how the mental life of a child changes dramatically, qualitatively, from one stage in which some tasks are impossible to a new stage in which those tasks become the everyday norm: each new stage represents a "super" intelligence from the viewpoint of the previous stage. There is an age at which the child conceives little more than herself and her parents. That child's brain just cannot conceive that there are other people and that people live on a planet and that the planet contains animals, trees, seas, mountains, etc; that you have to study and work; not to mention the mind-boggling affairs of sex and where children come from; and that some day you will die. All of this emerges later in life, each stage unlocking a new dimension of understanding. (And i wonder if there is an end to this process: if we lived to be 200 years old in good health, what would be our understanding?) My intelligence is "super" compared to the intelligence i had as a little child.

At the same time try learning languages or any other skills at the speed that children learn them. Children can do things with their minds that adults cannot do anymore: sometimes you feel that you cannot understand what their minds do, that they are little monsters. Children are superhuman too, as Alison Gopnik argues in "The Philosophical Baby" (2009). One wonders what we could achieve if we remained children all our lives (more about this later).

From Non-human to Super-human Intelligence

Non-human intelligence exists all around us, so then the question becomes: what would make a non-human intelligence also "superhuman"? i have not seen a definition yet, so everything that follows is vague in nature because i literally don't know what "superhuman" is supposed to mean (as opposed to simply "non-human"). I am not even sure whether superhuman intelligence

requires human intelligence first, or can human-level intelligence be skipped on the way to superhuman intelligence?

If human intelligence is not a required first step, then i'm not sure what the required first step is, and the discussion from vague becomes completely abstract.

If human intelligence is a required first step, isn't it easier to just build an electronic replica of a human brain rather than arguing about Artificial Intelligence? If we replaced each and every neuron in your brain with an electronic chip, i am not sure that you would still be "you" but your brain should still yield a form of human intelligence, shouldn't it?

Unfortunately, we are pretty far from implementing that full replica of your brain (note that i keep saying "your" and not "mine"). It is a bit discouraging that the smallest known brain, the brain of the roundworm (300 neurons connected by a few thousand synapses) is still smarter than the smartest neural network ever built.

If you think that this electronic replica of your brain is not as smart as you, then you are implying that the very "stuff" of which the brain is made is important in itself; but then machine intelligence is impossible.

I also wonder whether consciousness is necessary in order to be "intelligent". Does "machine intelligence" require all of you, including the mysterious inscrutable silent existence that populates your skull, the vast unexplored land of unspoken thoughts and feelings that constitutes "you"? What i hear when i listen to you is just a tiny fraction of what you have been and are thinking. What i see when i watch you is just a tiny fraction of what you thought of doing, dreamed of doing, and plan doing. When i hear a robot talk, that is the one and only thing that it "wants" to say. When i see it move, that is the one and only thing that it wants to do.

An electronic replica of your brain might or might not be conscious. You believe one or the other depending on your belief on where consciousness comes from. And that becomes a very long discussion. My multi-volume book "Thinking about Thought" is a survey of just the most influential viewpoints on consciousness.

And, to wit, “consciousness” for an information-processing machine could be something different from the consciousness of an energy-proceeding being like us. Our qualia (conscious feelings) measure energy levels: light, sound, etc. If information-processing machines ever develop qualia, it would make sense that those qualia be about information levels; not qualia related to physical life, but qualia related to “virtual” life in the universe of information

We don't really know how to build conscious beings (because we have no clue where consciousness comes from) and not even how to find out if something is conscious. However, i doubt i would call "superhuman" something that is less conscious than me. If you believe that consciousness is required for machine intelligence, you can stop reading here because nobody has a clue how to manufacture that.

The Accelerated Evolution of Machines

In all cases of rapid progress in the functionalities of a machine it is tempting to say that the machine achieved in a few years what took humans millions of years of evolution to achieve. The argument goes like this: "Yes, it took years to build a machine that recognizes a cat, but how long did it take evolution to create a living being that recognizes cats?"

My response is that any human-made technology is indirectly using the millions of years of evolution that it took to evolve its creator. No human being, no machine. Therefore it is incorrect to claim that the machine came out of nowhere: it came out of millions of years of evolution, just like my nose. The machine that is now so much better than previous models of a few years ago did NOT evolve: WE evolved it (and continue to evolve it).

There is no machine that has created another machine that is superior. WE create a better machine.

We are capable of doing that because those millions of years of evolution equipped us with some skills (that the machine does NOT have). If humans gets extinct tomorrow morning, the evolution of machines ends. Right now this is true of all technologies. If all humans die, all our technologies die with us (until a new form of intelligent life arises from millions of years of evolution and starts rebuilding all those watches, bikes, coffee makers, dishwashers, airplanes and computers). Hence, technically speaking, there has been no evolution of technology.

This is yet another case in which we are applying an attribute invented for one category of things to a different category: the category of living beings evolve, the category of machines does something else, which we call "evolve" by recycling a word that actually has a different meaning. It would be more appropriate to say that a technology "has been evolved" rather than "evolved": computers have been evolved rapidly (by humans) since their invention.

Technologies don't evolve (as of today): we make them evolve.

The day we have machines that survive without human intervention and that build other machines without human intervention, we can apply the word "evolve" to those machines.

As far as i know those machines don't exist yet, which means that there has been zero evolution in machine intelligence so far.

The machine is not intelligent, the engineer who designed it is. That engineer is the product of millions of years of evolution, the machine is a by-product of that engineer's millions of years of evolution.

(See the appendix for a provocative counter-argument).

Human Obsolescence

Several contemporary thinkers fear that us (humans) are becoming obsolete because machines will soon take our place. Jack Good wrote in "Speculations Concerning the First Ultraintelligent Machine" (1965): "the first ultraintelligent machine is

the last invention that man need ever make". Hans Moravec in "Mind Children" (1988): "robots will eventually succeed us: humans clearly face extinction". Erik Brynjolfsson and Andrew McAfee have analyzed the problem in "Race Against the Machine" (2012). Actually, this idea has been repeated often since the invention of (among other things) the typewriter and the assembly line.

In order to understand what we are talking about we need to define what is "us". Assembly lines, typewriters, computers, search engines and whatever comes next have replaced jobs that have to do with material life. I could simply say that they have replaced "jobs". They have not replaced "people". They replaced their jobs. Therefore what went obsolete has been jobs, not people, and what is becoming obsolete is jobs, not people. Humans, to me, are biological organisms who (and not "that") write novels, compose music, make films, play soccer, ride the Tour de France, discover scientific theories, hike on mountains and dine at fancy restaurants. Which of these activities is becoming obsolete because machines are doing them better?

Machines are certainly good at processing big data at lightning speed. Fine. We are rapidly becoming obsolete at doing that. In fact, we've never done that. Very few humans spent their time analyzing big data. The vast majority of people are perfectly content with small data: the price of gasoline, the name of the president, the standings in the soccer league, the change in my pocket, the amount of my electricity bill, my address, etc. Humans have mostly been annoyed by big data. That was, in fact, a motivation to invent a machine that would take care of big data. The motivation to invent a machine that rides the Tour de France is minimal because we actually enjoy watching (human) riders sweat on those steep mountain roads, and many of us enjoy emulating them on the hills behind our home. Big data? Soon we will have a generation that cannot even do arithmetic.

So we can agree that what is becoming obsolete is not "us" but our current jobs. That has been the case since the invention of the first farm (that made obsolete the prehistoric gatherers) and, in fact,

since the invention of the wheel (that probably made obsolete many who were making a living carrying goods on their backs).

On the other hand, every machine invented so far has become obsolete very (very) rapidly. I suspect the future will be more like the past.

Historical Intermezzo: Jobs in the 2000s

During the Great Depression of 2008-2011 both analysts and ordinary families were looking for culprits to blame for the high rate of unemployment, and automation became a popular one in the developed world. Automation is indeed responsible for making many jobs obsolete in the 2000s, but it is not the only culprit.

The first and major one is the end of the Cold War. Before 1991 the economies that really mattered were a handful (USA, Japan, Western Europe). Since 1991 the number of competitors for the industrialized countries has skyrocketed, and they are becoming better and better. Technology might have "stolen" some jobs, but that factor pales by comparison with the millions of jobs that were exported to Asia. In fact, if one considers the totality of the world, an incredible number of jobs have been created precisely during the period in which critics argue that millions of jobs have been lost. If Kansas loses one thousand jobs but California creates two thousand, we consider it an increase in employment. These critics make the mistake of using the old nation-based logic for the globalized world. When counting jobs lost or created during the last twenty years, one needs to consider the entire interconnected economic system. Talking about the employment data for the USA but saying nothing about the employment data (over the same period) of China, India, Mexico and so forth is distorting the picture. If General Motors lays off one thousand employees in Michigan but hires two thousand in China, it is not correct to simply conclude that "one thousand jobs have been lost". If the car industry in the USA loses ten thousand jobs but the car industry in China gains twenty thousand, it is not correct to simply conclude that ten thousand jobs

have been lost by the car industry. In these cases jobs have actually been created.

That was precisely the case: millions of jobs have been created by the USA in the rest of the world while millions were lost at home. The big driver was not automation but, on the contrary, cheap labor.

Then there are sociopolitical factors. Unemployment is high in Western Europe, especially among young people, not because of technology but because of rigid labor laws and government debt. A company that cannot lay off workers is reluctant to hire any. A government that is indebted cannot pump money into the economy. This is a widespread problem in the Western economies of 2013, but not in the economies of the rest of the world. It has to do with politics, not with automation.

Germany is as technologically advanced as the USA. All sorts of jobs have been fully automated. And, still, in Germany the average hourly pay has risen five times faster between 1985 and 2012 than in the USA. This has little to do with automation: it has to do with the laws of the country. Hedrick Smith's "Who Stole the American Dream?" (2012) lays the blame on many factors, but not on automation.

In 1953 Taiichi Ohno invented "lean manufacturing" at Japan's Toyota, possibly the most important revolution in manufacturing since Ford's assembly line. Nonetheless, Japan created millions of jobs in manufacturing; and, in fact, Toyota went on to become the largest employer in the world of car-manufacturing jobs. Even throughout its two "lost decades" Japan continued to post very low unemployment.

Another major factor that accounts for massive losses of jobs in the developed world is the management science that emerged in the 1920s in the USA. That science is the main reason that today companies don't need as many employees as comparable companies employed a century ago. Each generation of companies has been "slimmer" than the previous generation. As those management techniques get codified and applied across the board, companies become more efficient at manufacturing (across the

world) and selling (using the most efficient channels) and at predicting business cycles. All of this results in fewer employees not because of automation but because of optimization.

Additionally, in the new century the USA has deliberately restricted immigration to the point that thousands of brains are sent back to their home countries even after they graduated in the USA. This is a number that is virtually impossible to estimate, but, in a free market like the USA that encourages innovation and startups, jobs are mostly created via innovation, and innovation comes from the best brains, which account for a tiny percentage of the population. Whenever the USA sends back or refuses to accept a foreign brain that may have become one of those creators of innovation, the USA is de facto erasing thousands of future jobs. Those brains are trapped in places where the system does not encourage the startup-kind of innovation (like Western Europe and China) or where capital is not as readily available (like India). They are wasted in a way that equivalent brains were not wasted in the days when immigration into the USA was much easier, up until the generation of Yahoo, eBay and Google. The "Kauffman Thoughtbook 2009" by the Ewing Marion Kauffman Foundation contains a study that foreign-born entrepreneurs ran 24% of the technology businesses started between 1980 and 1998 (in Silicon Valley a staggering 52%). In 2005 these companies generated \$52 billion in revenue and employed 450,000 workers. In 2011 a report from the Partnership for a New American Economy found that 18% of the Fortune 500 companies of 2010 were founded by immigrants. These companies had combined revenues of \$1.7 trillion and employed millions of workers. (If one includes the Fortune 500 companies founded by children of immigrants, the combined revenues were \$4.2 trillion in 2010, greater than the GDP of any other country in the world except China and Japan).

Technology is certainly a factor, but it can go either way. Take, for example, energy. This is the age of energy. Energy has always been important for economic activity but never like in this century. The cost and availability of energy are major factors to determine

growth rates and therefore employment. The higher the cost of energy, the lower the amount of goods that can be produced, the lower the number of people that we employ. If forecasts by international agencies are correct, the coming energy boom in the USA (see the International Energy Agency's "World Energy Outlook" of 2012) might create millions of jobs, both directly and indirectly. That energy boom is due to new technology.

When the digital communication and automation technologies first became widespread, it was widely forecasted that a) people would start working from home and b) people would not need to work as much. What I have witnessed is the exact opposite: virtually every company in Silicon Valley requires people to show up at work a lot more than they did in the 1980s, and today virtually everybody is "plugged in" all the time. I have friends who check their email nonstop while we are driving to the mountains and even while we are hiking. The digital communication and automation technologies have not resulted in machines replacing these engineers but in these engineers being able to work all the time from everywhere, and their companies require it. Those technologies have resulted in people working a lot more. (The willingness of people to work more hours for free is another rarely mentioned factor that is contributing to higher unemployment).

Unemployment cannot be explained simply by looking at the effects of technology. Technology is one of many factors and, so far, not the main one. There have been periods of rapid technological progress that have actually resulted in very low unemployment (i.e. lots of jobs), most recently in the 1990s when e-commerce was introduced, despite the fact that the digital camera killed the photographer's shop, Amazon killed the bookstore, the cell phone killed the land lines and Craigslist killed the local newspaper.

Historically, in fact, technology has created jobs while it was destroying old jobs, and the new jobs have typically been better-paying and safer than the old ones. Not many people dream of returning to the old days when agriculture was fully manual and

millions of people were working in terrible conditions in the fields. Today a few machines can water, seed, rototill and reap a large field. Those jobs don't exist anymore, but many jobs have been created in manufacturing sectors for designing and building those machines. Few peasants in the world would like their children to grow up to be peasants instead of mechanical engineers. Ditto for computers that replaced typewriters and typewriters that replaced pens and pens that replaced human memory.

It is true that the largest companies of the 21st century are much smaller than the largest companies of the 20th century. However, the world's 4,000 largest companies spend more than 50% of their revenues on their suppliers and a much smaller percentage on their people (as little as 12% according to some studies). Apple may not be as big as IBM was when it was comparable in power, but Apple is the reason that hundreds of thousands of people have jobs in companies that make the parts that Apple products use.

I would be much more worried about the "gift economy": the fact that millions of people are so eager to contribute content and services for free on the Internet. For example, the reason that journalists are losing their jobs has little to do with the automation in their departments and a lot to do with the millions of people who provide content for free on the Internet.

If one takes into account the real causes of the high unemployment rate in the USA and Europe, we reach different conclusions about the impacts that robots (automation in general) will have. In the USA robots are likely to bring back jobs. The whole point of exporting jobs to Asia was to benefit from the lower wages of Asian countries; but a robot that works for free 24 hours a day 7 days a week beats even the exploited workers of communist China. As they become more affordable, these "robots" (automation in general) will displace Chinese workers, not Michigan workers. The short-term impact will be to make outsourcing of manufacturing an obsolete concept. The large corporations that shifted thousands of jobs to Asia will bring them back. In the mid term, if this works out well, a secondary effect will be to put Chinese products out of the

market and create a manufacturing boom in the USA: not only old jobs will come back but a lot of new jobs will be created. In the long term robots might create new kinds of jobs that today we cannot foresee. Not many people in 1946 realized that millions of software engineers would be required by the computer industry in 2013. My guess is that millions of "robot engineers" will be required in a heavily robotic future. Those engineers will not be as "smart" as their robots at whatever task for which those robots were designed just like today's software engineers are not as fast as the programs they create. And my guess is that robots will become obsolete too at some point, replaced by something else that today doesn't even have a name. If i had to bet, i would bet that robots (intelligent machines in general) will become obsolete way before humans become obsolete.

The real revolution in employment is coming from a different direction: the "sharing economy". Companies such as Airbnb, that matches people who own rooms and people looking for rooms to rent, and Uber, that matches drivers who own a car and people looking for a ride in a car, have introduced a revolutionary paradigm in the job market: let people monetize under-utilized assets. This concept will soon be applied in dozens of different fields, allowing ordinary people to find ordinary customers for their ordinary assets; or, in other words, to supply labor and skills on demand. Before the industrial revolution most jobs were in the countryside but urban industry existed and it consisted mainly of artisan shops. The artisans would occasionally travel to a regional market, but mostly it was the customer who looked for the artisan, not viceversa. Cities like Firenze (Florence) had streets devoted to specific crafts, so that a customer could easily find where all the artisans offering a certain product were located. Then came the age of the factory and of transportation, and industrialization created the "firm" employing thousands of workers organized in some kind of hierarchy. Having a job came to mean something else: being employed. Eventually society started counting "unemployed" people, i.e. people who would like to work for an employer but no employer wants their time

or skills. The smartphone and the Internet are enabling a return of sorts to the model of the artisan era. Anybody can offer their time and skills to anybody who wants them. The "firm" is simply the intermediary that allows customers to find the modern equivalent of the artisan.

In a sense, the "firm" (such as Uber or Airbnb) plays the role that the artisan street used to play in Firenze. Everybody who has time and/or skills to offer can now become a "self-employed" person. And that "self-employed" person can work when she wants, not necessarily from 8 to 5. There is no need for an office and for hiring contracts.

The traditional firm has a workforce that needs to be fully employed all the time, and sometimes the firm has to lay off workers and sometimes has to hire some, according to complicated strategic calculations.

In the sharing economy, no such thing exists: the firm is replaced by a community of skilled workers who take the jobs they want to take when they want to take them if the customer wants them to take them. In a sense, people can be fired and hired on the fly.

Of course, this means that "good jobs" will no longer be judged based on job promotions, salary increases and benefits. They will be based on customer demand (which in theory is what drives company's revenues which in turn drives job promotions, salary increases and benefits).

The unemployed person who finds it difficult to find a job in a firm is someone whose skill is not desired by any firm, but this does not mean that his skills are not desired by any customer. The firm introduced a huge interface between customer and worker. When there is a need for your skill, you have to hope that a manager learns of your skills, usually represented by a resume that you submitted to the human resources department, and hope that the financial officer will approve the hiring. The simple match-making between a customer who wants a service and the skilled worker who can provide that service gets complicated by the nature of the firm with its hierarchical structure and its system of checks and

balances (not to mention internal politics and managerial incompetence). It would obviously be easier to let the customer deal directly with the skilled worker who can offer the required service.

Until the 2000s the problem was that the customer had no easy way to access skilled workers other than through the "yellow pages", i.e. the firms. Internet-based sharing systems remove the layers of intermediaries except one (the match-making platform, which basically provides the economy of scale). In fact, these platform turn the model upside down: instead of a worker looking for employment in a firm that is looking for customers, the new model views customers as looking for workers. Not only does this model bypass the slow and dumb firm, but it also allows you to monetize assets that you own and you never perceived as assets. A car is an asset. You use it to go to work and to go on vacation, but, when it is parked in the garage, it is an under-utilized asset.

Marketing used to be a scientific process to shovel a new product down the throat of reluctant consumers: it now becomes a simple algorithm allowing customers to pick their skilled workers, an algorithm that basically combines the technology of online dating (match making), of auctions (bidding) and of consumer rating (that basically replaces the traditional "performance appraisal" prescribed in the traditional firm).

Of course, the downside of this new economy is that the worker has none of the protections that she had in the old economy: no security that tomorrow she will make money, no corporate pension plan, etc; and she is in charge of training herself to keep herself competitive in her business. The responsibility for a worker's future was mostly offloaded to the firm. In the sharing economy that responsibility shifts entirely to the worker herself.

The new proletariat is self-employed, and, basically, each member of the proletariat is actually a micro-capitalist; the price to pay is that the worker will have to shoulder the same responsibilities that traditionally have fallen into the realm of firm management.

Futurists have a unique way to completely miss the scientific revolutions that really matter.

Marketing and Fashion

Back to the topic of accelerating progress: what is truly accelerating at exponential speed is fashion. This is another point where many futurists and high-tech bloggers confuse a sociopolitical phenomenon with a technological phenomenon.

What we are actually witnessing in many fields is a regression in quality. This is largely due to the level of sophistication reached by marketing techniques. Marketing is a scary human invention: it often consists in erasing the memory of good things so that people will buy bad things. There would be no market for new films or books if everybody knew about the thousands of good films and books of the past: people would spend their entire life watching and reading the (far superior) classics instead of the new films and books, most of which are mediocre at best. In order to have people watch a new film or read a new book, the marketing strategists have to make sure that people will never know about old films and books. It is often ignorance that makes people think they just witnessed "progress" in any publicized event. Often we call "progress" the fact that a company is getting rich by selling poor quality. The "progress" lies in the marketing, not in the goods. The acceleration of complexity is in reality an acceleration of low quality.

We may or may not live in the age of machines, but we certainly live in the age of marketing. If we did not invent anything, absolutely anything, there would still be frantic change. Today change is largely driven by marketing. The industry desperately needs consumers to go out and keep buying newer models of everything. We mostly buy things we don't need. The younger generation is always more likely to be duped by marketing and soon the older generations find themselves unable to communicate with young people unless they too buy the same things. Sure: many of them are convenient and soon come to be perceived as "necessities"; but the truth is that humans have lived well (sometimes better) for millennia without those "necessities". The idea that an mp3 file is better than a compact disc which is better than a vinyl record is just

that: an idea, and mainly a marketing idea. The idea that a streamed movie is better than a DVD which is better than a VHS tape is just that: an idea, and mainly a marketing idea. We live in the age of consumerism, of rapid and continuous change in products, mostly unnecessary ones.

What is truly accelerating is the ability of marketing strategies to create the need for new products. Therefore, yes, our world is changing more rapidly than ever; not because we are surrounded by better machines but because we are surrounded by better snake-oil peddlers (and dumber consumers).

"The computer industry is the only industry that is more fashion-driven than women's fashion" (Larry Ellison, founder and chairman of Oracle).

Sometimes we are confusing progress in management, manufacturing and marketing (that accounts for 90 percent of the "accelerating progress" that we experience) with progress in machine intelligence (that is still at the "Press 1 for English" level).

Technological progress is, in turn, largely driven by its ability to increase sales. Therefore it is not surprising that the big success stories of the World-wide Web (Yahoo, Google, Facebook, etc) are the ones that managed to turn web traffic into advertising revenues. We are turning search engines, social media and just about every website into the equivalent of the billboards that dot city streets and highways. It is advertising revenues, not the aim of creating intelligent machines, that is driving progress on the Internet. In a sense, Internet technology was initially driven by the military establishment, that wanted to protect the USA from a nuclear strike, then by a utopian community of scientists that wanted to share knowledge, then by corporations that wanted to profit from e-commerce, and now by managers of advertising campaigns who want to capture as large an audience as possible. Whether this helps accelerate progress and in which direction is, at best, not clear.

When Vance Packard wrote his pamphlet "The Hidden Persuaders" (1957) on the advertising industry (on how the media

can create the illusory need for unnecessary goods), he had literally seen nothing yet.

"The best minds of my generation are thinking about how to make people click ads" (former Facebook research scientist Jeff Hammerbacher in 2012).

And, to be fair, the best minds of his generation are not only used to make people click on ads but also to create ever more sophisticated programs of mass surveillance (as revealed in 2013 by National Security Agency analyst Edward Snowden).

The Audience of the Singularity

I organize many events in the San Francisco Bay Area. I am always frustrated that so few young people show up. I routinely attend technical and scientific talks at prestigious organizations like Stanford University, the Computer History Museum and Xerox PARC. They are free and frequently feature top-notch speakers. At least half of the audience is consistently made of gray-haired people. They constitute a tiny percentage of the active workforce, especially in high tech. Therefore, it is mostly older inactive engineers who hear the distinguished researchers talk about the state of high technology. The percentage of young people who learn about the state of technology is relatively small. The reason why younger people don't come in proportional numbers to educational events is simple: they are busy at work or studying, and, why not, they might be fed up after so many years of college and just want to party in the evening. Younger people are therefore more likely to get their technology news from attending yearly conferences and trade shows and (on a daily basis) from reading popular bloggers. What they get is, in other words, press releases. (Don't even try to convince me that your favorite tech blogger is competent and reliable: he is just a flywheel in a highly efficient system to distribute press releases by high-tech companies, and mostly product announcements, with little or no knowledge of the science behind technology and little or no contacts in the labs that

produced that science before some startup turned it into a popular gadget). Therefore young technology buffs are more likely to welcome enthusiastically the news that some startup has introduced a new battery that will last centuries because they never heard the head of ARPA-E talk about the state of battery technology (the startup that made that announcement is simply looking for funds from venture capitalists and needs to create buzz around its business plan). They are also the ones who tend to believe that Artificial Intelligence has built incredibly smart machines and that the Singularity is coming soon.

That is half of the audience that absorbs enthusiastically any news about machine intelligence. The other half is the one that i compared to religiously devout people who simply have an optimistic view of all these press releases (i obviously have a much more skeptical view).

In Defense of Progress

Enough bashing machines. It is true that i am basically living in the world of my parents, but the one difference (the computer) cannot be considered as just one item: it is many applications in one. What can the dishwasher do other than wash dishes? The computer, instead, can do a lot of things, from delivering mail to displaying pictures.

That was, in fact, the whole point of the Turing machine: a universal problem solver. Little did he know that its applications would range from phone conversations to social networking.

There has been little progress in the physical world but a lot in the virtual world created by computers. Just witness the explosion of online services of the late 1990s and of smartphone applications in the last few years.

Perhaps even more importantly, the law of entropy does not apply to software: everything in this universe is bound to decay and die because of the second law of Thermodynamics (that entropy can never decrease). That does not apply to software. Software will

never decay. Software can create worlds in which the second law of Thermodynamics does not apply: software things never age, never decay, never die. (Unfortunately, software needs hardware to run, and that hardware does decay).

One could argue that Artificial Intelligence has failed to deliver, but "Augmented Intelligence" has been successful beyond the hopes of its founding fathers. In the 1960s in Silicon Valley there were two schools of thought. One, usually associated with John McCarthy's Stanford Artificial Intelligence Lab (SAIL), claimed that machines would soon replace humans. The other one, mainly associated with Doug Engelbart at the nearby Stanford Research Institute (now SRI Intl), argued that machines would "augment" human intelligence rather than replace it. Engelbart's school went on to invent the graphic user interface, the personal computer and the Internet, among many other things, all the way down to virtual personal assistants like SIRI; all things that "augmented" human intelligence. This program did not necessarily increase human intelligence and it did not create a non-human intelligence: the combination of human intelligence plus these devices can achieve "more" than human intelligence can alone. These are tools that may or may not increase the intelligence of the user, just like good books, good music, hiking, travel, etc.

Don't get me wrong: i have a lot of respect for all those algorithms that combine the best of Computational Mathematics and for all the engineering that lies behind today's multi-functional gadgets. For example, it must be terribly difficult for search engines to keep up with the exponential growth of user-provided content. The ranking algorithm has to become exponentially smarter in order for the search engine to keep providing relevant answers. It's something that the user doesn't see (unlike, say, a new button on the microwave oven), but it's something vital lest the search engine becomes yet another useless website and the World-wide Web becomes unsearchable (i.e. it becomes the "World-wide Mess").

Intermezzo: You Are a Gadget

The combination of phones, computers and networks has put each individual in touch with a great number of other individuals, more than at any time in history: humankind at your fingertips. This is certainly useful for scholars who need feedback from readers. This is certainly lucrative for businesses that want to reach as many consumers as possible with their adverts. But do ordinary people really benefit from being connected to thousands of people, and soon millions? What happens to solitude, meditation, to "thinking" in general (whether scientific thinking or personal recollection) when we are constantly interacting with a multitude of minds (only some of which really care)?

A general rule is that you "are" the people with whom you interact, because they influence who you become. In the old days those were friends, relatives, neighbors and coworkers. Now they are strangers spread all over the world (and old acquaintances with whom you only share distant memories). Do you really want to be "them" rather than being yourself?

If you surround yourself with philosophers, you are likely to become a philosopher, even if only an amateur one. If you surround yourself with book readers, you are likely to read a lot of books. And so on. So what is likely to happen to you if you surround yourself with gadgets that mediate your interaction with people and with the world at large?

It is infinitely easier to produce/accumulate information than to understand it and make others understand it.

Task-specific vs General-purpose Intelligence

Before analyzing what it will take (and how long it will take) to get machine intelligence we need to define what we are talking about.

A man wearing a suit and tie walks out from a revolving hotel door dragging his rolling suitcase. Later another man, wearing a uniform and gloves, walks out of the side door dragging a garbage can. It is obvious even to the dumbest human being that one is a guest of the hotel and the other one is a janitor. Do we require from a machine

this simple kind of understanding ordinary situations in order for it to qualify as "intelligent"? Or is it irrelevant, just like matching the nightingale's song is irrelevant in order to solve differential equations? If we require that kind of understanding, we push machine intelligence dramatically forward into the future: just figuring out that one is a suit and tie and one is a uniform is not trivial at all for a machine. It will take an enormous effort to achieve just this one task. There are millions of similarly "ordinary" tasks that humans routinely perform without even realizing that they are. There are millions of situations that we recognize in a split second.

Let us continue our thought experiment. Now we are in an underdeveloped country and the janitor is dragging not a garbage can which is an old broken suitcase full of garbage. He has turned an old suitcase into his garbage can. We would probably just smile seeing such a scene; but imagine how hard it is for a machine to realize what is going on. Even if the machine is capable of telling that someone dragging a suitcase is a hotel guest, the machine now has to understand that a broken suitcase carried by a person in a janitor's uniform does not qualify as a suitcase.

There are millions of variants on each of those millions of situations that we effortlessly understand, but that are increasingly trickier for a machine.

The way that today's A.I. scientists would go about it is to create one specific program for each of the millions of situations, and then millions of their variants. Given enough engineers, time and processors, this is feasible. Whenever a critic like me asks "but can your machine do this too?", today's A.I. scientists rush out to create a new machine that can do it. "But can your machine also do this other thing?" The A.I. scientists rush out to create another machine. And so forth.

Given enough engineers, time and processors, it is indeed possible to create a million of a million machines that can do everything we naturally do.

This is equivalent to the idea that the Web plus a search engine can answer any question: someone, sooner or later, will post the

answer on the Web, and the search engine will find it. Billions of Web users are providing all the answers to all the possible questions. The search engine is not particularly intelligent in any field but can find the answer to questions in all fields.

I doubt that this is the way in which my mind works (or any animal's mind works), but, yes, those millions of machines will be functionally equivalent to my mind. In fact, they will be better than my mind because they will be able to understand all situations that all people understand, not just the ones that i understand, just like the Web will eventually contain the answers to all questions that all humans can answer, not only the answers that i know.

Incidentally, the effect on the economy will be to create millions of jobs because those millions of machines will need to be designed, tested, stored, marketed, sold, and, last but not least, repaired. (I am a bit more skeptic that automation will benefit small businesses over the multinationals, one of the many promises that the Internet failed to deliver: Amazon killed thousands of bookstores around the world and something similar will probably happen in the robotic world).

Intelligent Behavior from Structured Environments

When you need to catch a bus in an underdeveloped country, you don't know what time it will arrive nor how much you will be charged for the ticket. In fact you don't even know how it will look like (it could be a generic truck or a minivan) and where it will stop. Once on board, you tell the driver where you want to get off and hope that he will remember. If she is in a good mood, after a little amusing chat, she might even take a little detour to drop you right in front of your hotel. When you take a bus in a developed country, there is an official bus stop (the bus won't stop if you are 20 meters before or after it), the bus is clearly recognizable and marked with the destination and won't take any detour for any reason, the driver is not allowed to chat with the passengers (sometimes she is physically enclosed in a glass cage), the ticket must be bought with

exact change at a ticket vending machine (and sometimes validated inside at another machine). There is a door to be used to exit, and you know when to exit because the name of the bus stop is displayed on a screen. In the case of most trains and many buses, you also get an assigned seat (you can't just sit anywhere).

It is easy to build a robot that can ride a bus in a developed country, much more difficult to build a robot that can ride a bus in an underdeveloped country. What makes it easy or difficult is the environment in which it has to operate: the more structured the environment, the easier for the robot. A structured environment requires less "thinking": just follow the rules and you'll make it. However, what really "makes it" is not you: it's you plus the structured environment. That's the key difference: operating in a chaotic, unpredictable situation is not the same thing as operating in a highly structured environment. The environment makes a huge difference. It is easy to build a machine that has to operate in a highly structured environment, just like it is easy for a bullet train to ride at 300 km/hour on rails.

We structure the chaos of nature because it makes it easier to survive and thrive in it. Humans have been spectacularly successful at structuring their environment so that it obeys simple, predictable rules. This way we don't need to "think" too much: the environment will take us where we want to go. We know that we can find food at the supermarket and a train at the train station. In other words, the environment makes us a little more stupid but allows anybody to achieve tasks that would otherwise be difficult and dangerous, i.e. that would require a lot of intelligence. When the environment fails us, we get upset because now we have to think, we have to find a solution to an unstructured problem. If you are in Paris and the metro is on strike and it is impossible to get a taxi, how do you get to your appointment in time? Believe it or not, most Parisians manage. Of course, most tourists from the USA don't. If there is no traffic light and cars don't stop for pedestrians and traffic is absolutely horrible, how do you cross a wide boulevard? Believe it or not, Iranians do it all the time. Needless to say, most Western

tourists spend hours trying to figure it out. It is certainly very impressive how well humans structure a universe that is chaotic. The more we structure it, the easier for extremely dumb people and machines to survive and thrive in it.

The claims of the robotic industry are often related to structured environments, not to their robots. It is relatively easy to build an autonomous car that rides on a highway with clearly marked lanes, clearly marked exits, ordered traffic, and maps that detail everything that is going to happen. It is much more difficult (orders of magnitude more difficult) to build an autonomous car that can drive through Tehran or Lagos (this is a compliment to Iranian and Nigerian drivers, not an insult). Whoever claims that a computer is driving a car is distorting the facts: it is not the computer that is driving the car but the environment that has been structured so that any inexperienced and not particularly intelligent driver, and even a computer, can drive a car. The computer cannot drive a car in the traffic of Lagos or Tehran. It will if and when the streets of Lagos and Tehran become as well structured as the streets of California, if and when Iranian and Nigerian drivers are forced to obey strict traffic rules. Saying that the on-board computer is steering the driverless car is like saying that the locomotive knows in which direction to take the train: the locomotive is simply constrained by the rails to take the correct direction.

In fact, the feats of some A.I. systems are rarely repeated in other laboratories: those systems wouldn't work in other laboratories. They need the structured environment of their native laboratory in order to function. What really "functions" is the environment in which they are dropped.

I recently had to exchange the equivalent of \$3.00 in a local currency while leaving a Western country at its capital's airport. The procedure was silly beyond belief. I had to produce passport, boarding pass and receipt of previous money exchanges before getting my money, a lengthy operation for just three dollars. On the contrary at the border between Haiti and Dominican Republic, a wildly chaotic place with taxi drivers, fruit vendors and police

officers yelling at each other and at everybody passing by, there was a mob of money changers chasing the travelers. I had to guess which ones were honest money changers from the scammers, then bargain the exchange rate, then make sure that the money was good while all the time protecting my wallet from pickpockets. It wouldn't be difficult to build a robot that can exchange money at the airport of a Western capital, but much more difficult (orders of magnitude more difficult) to build one that can exchange money while walking from the immigration post of Haiti to the immigration post of the Dominican Republic.

The point, again, is: the more structured the environment the easier to build a machine that operates in it. What really "does it" is not the machine: it's the structured environment. What has made so many machines possible is not a better A.I. technology, but simply more structured environments. It's the rules and regulations that allow the machine to operate.

You can't call an automatic phone system and just explain your problem. You have to press 1 for English, 1 for customer support, 3 for your location, 2 for your kind of problem and 4 and 7 and so forth. What allows the machine to perform its job, and to replace the human operator, is that you (the human being) have removed the human aspect from the interaction and behave like a machine in a mechanic world. It is not the machine that behaves like a human being in a human world.

The fundamental thing that a self-driving car must be able to do is, of course, to stop at a gas station when it runs out of gasoline. Can these cars autonomously enter a gas station, stop in front of a pump, slide a credit card in the payment slot, pull out the hose and pour gasoline in the tank? Of course, not. What needs to be done is to create the appropriate structured environment for the driverless car (or, better, for some sensors on board of the car) so that the car will NOT need to behave like an intelligent being. The gas station, the gas pump and the payment used by the driverless car will look very different from the one used so far by human drivers.

Incidentally, most of those rules and regulations that create a highly structured environment (favorable to automata) were originally introduced in order to reduce costs. Employing machines has been the next logical step in cost reduction. The machine is one step in an ongoing process of cost reduction, productivity increase, etc. The goal was not to create superhuman intelligence, just to increase profits.

The more unstructured the environment is the more unlikely that a machine can replace the human. Unfortunately, in 2013 one very unstructured environment is health care. Medical records are kept on paper files, and doctor's notes are notoriously impossible to read. There is very little that a machine can do in that environment. The way to introduce "intelligent" machines in that environment is, first of all, to structure all that information. When it is "digitized", it means that it has been structured. At that point any human being, even with little or no knowledge of medical practice, can do something intelligent in that environment. And even a machine can.

Think of your favorite sandwich chain. You know exactly what kind of questions they will ask you. There is a well-structured process by which your sandwich will be made. The moment robots become cheap enough they will certainly take over the jobs of the kids who prepare your sandwich today. It is not a matter of "intelligence" (the intelligence of today's robots is already more than enough) but of cost: today a teenager is cheaper than a robot. The whole point of structuring the sandwich-making process was to allow inexperienced and unskilled workers (read: underpaid) to perform the task once reserved to skilled experienced chefs.

The truth is that we do not automate jobs as they are. First, we dehumanize the job, turning it into a mechanical sequence of steps. Then we use a machine to automate what is left of that job. For example, my friend Steve Kaufman, a pediatrician all his life, realized that his skills were less and less necessary: a nurse practitioner can fill all the forms and click on all the computer buttons that are required when seeing a patient; the doctor, who is increasingly required to type on a keyboard, may not even make

eye contact with the patient. This has the beneficial effect of reducing the number of days that the average patient spends at a hospital, but it erases the kind of bonding between doctor and patient that was common in the “unstructured” world. When the last vestiges of humanity will have been removed from the job of the doctor, it will be relatively easy to automate the doctor’s job. But that is not what Steve was doing. As Steve pointed out to me, if you don’t bond with an asthmatic patient, you may never realize that he is suicidal: you will cure his asthma, but he will commit suicide; and the machine will archive the case as a success.

Structured environments are also relying on ever stricter rules. My favorite example is the boarding procedure at an airport, where we are treated like cattle from check-in to the gate, with a brief interval during which we are treated like a walking credit card with money in the bank that airport shops desperately try to get. Other than the credit card thing, we are basically building the kind of bureaucratic state pioneered by the Soviet Union.

Politics aside, there is a fundamental paradox underlying the ongoing structuring of society. What is profoundly human (and actually shared by all forms of life) is the vagueness of language and behavior. What humans (and animals) can do relatively well, and do on a daily basis, and machines are not good at is to deal with ambiguity. Unfortunately, ambiguity is responsible for a lot of the miscommunication and chaos that complicate our life. What rules and regulations do is to remove ambiguity from society, and therefore to simplify individual life. As a side-effect, the more we structure human behavior by removing ambiguity, the more replicable it becomes.

Increasingly structured environments, routines and practices will eventually enable the automation of “cognitive” skills too. As political debates are becoming more and more structured, with a format agreed beforehand and a moderator that enforces it, and a restriction on the kind of questions that can be asked, and candidates who basically memorize press releases worded by their campaign staff, it is not difficult to imagine that sooner or later

someone will build a piece of software that can credibly replace a politician in a political debate; but that feat will owe more to the lack of real debate in these political debates than to greater rhetorical skills by the machine. On the other hand that software will be incapable of participating in a passionate conversation about a World Cup game with a group of rowdy soccer fans.

In concluding, it is the structured environment itself that is enabling and will enable the explosion of robotics and automated services. Most of the robots and phone-based services coming to the market now rely on relatively old technology. What has made them feasible and practical is that they can now operate in highly structured environments.

Think of yourself. You are now identified by numbers in so many different contexts: your passport number, your social security number, your street address, your telephone number, your insurance policy number, your bank account number, your credit card number, your driver license number, your car's plate number, your utility bill account number... It is a rarity when someone tries to identify me based on non-numeric features. And increasingly we depend on passwords to access our own information. The more we reduce the individual to a digital file the easier it gets to build "intelligent assistants" for that file... sorry, i meant "for that individual".

In a sense, humans are trying to build machines that think like humans while machines are already building humans who think like machines.

Intermezzo: Will Intelligent Machines Return to Chaotic Environments?

For yet another sociological take on the discussion, examine what "structured environment" really means. On the one hand, it means removing the chaotic and unpredictable (and often intractable) behavior of natural environments. On the other hand, it also means removing the chaotic and unpredictable (and often intractable)

behavior of human beings. The purpose of all the rules and regulations that come with a structured environment is to replace you (a messy human intelligence) with an avatar that is like you (in fact it shares your body and brain) without the quirks of human intelligence, and that lives in a highly-structured virtual world that mimics the natural world without all the quirks of the (wildly unstructured) natural world.

My thesis is that machines are not becoming particularly more intelligent, but, instead, it is humans who are structuring the environment and regulating behavior so that humans become more like machines and therefore machines can replace humans.

Then the question is what happens if machines become "intelligent". If "intelligent" means that machines will become what humans are before society turns them into rule-obeying machines, then, ironically, machines may acquire all the "baggage" that intelligent biological beings carry, i.e. the unpredictable, chaotic, anarchic behavior that any living being exhibits, i.e. precisely what the structured environment and rules and regulations aim at suppressing.

A more rigid society, in which rules and regulations enforce some behavior and prohibit some other behavior, leaving little to the imagination, is also likely to have another side effect. Noise is important in self-organizing systems such as the human society because it allows such systems to evolve. Under the right circumstances a self-organizing system disturbed by noise will self-organize at a higher level, reaching a new level of equilibrium but in some cases profoundly different from the original one. The more we remove "noise", ambiguity and unpredictability from human society, the more we reduce the likelihood that human society will evolve at all, let alone towards higher levels of organization. Intelligent machines might rediscover this law of non-equilibrium thermodynamics after humans have forgotten it.

It would be ironic if creating intelligent machines would turn machines into (messy) humans at the same time that we are turning humans into (disciplined) machines.

The Multiplication of Appliances and Artificial Intelligence by Enumeration

If we structure the world appropriately, it will be easy to build machines that can board planes, exchange money, take a bus, drive a car, cross a street and so on. Automated services have existed since at least the invention of the waterwheel. We even have machines that dispense money (ATMs), machines that wash clothes (washing machines), machines that control the temperature of a room (thermostats), and machines that control the speed of a car (cruise controls).

When we design robots, we are simply building more appliances. What we might witness soon is a multiplication of appliances, disguised and marketed as "robots" simply because it is becoming fashionable.

The ATM is more precise than a bank clerk (and works much longer hours) but we don't think of it as "intelligent". Ditto for the washing machines that is capable of all sorts of washing techniques. Hence, it is not enough to just take over one human chore for a machine to be deemed "intelligent".

Enthusiastic fans of automation predict that "soon" (how soon?) everything that humans can do will be done by machines; but they rarely explain what is the point of making machines for everything we do. Do we really want machines that fall asleep or urinate? There are very human functions that people don't normally associate with "intelligence". They just happen to be things that human bodies do. We swing arms when we walk, but we don't consider "swinging arms while walking" a necessary feature of intelligent beings. The moment we attempt to design an "intelligent" machine (or collection of machines) that can mimic the entire repertory of our "intelligent" functions we run into the enumeration problem: which function qualifies as "intelligent"? Typical human activities include: forgetting where we left the cell phone, eating fast food, watching stand-up comedy, catching a flue when attacked by viruses and, yes, frequently, urinating.

One can decide that there is a hierarchy of tasks, from "not intelligent at all" to "very intelligent", and that the latter are the ones that make the difference; which is fine, but then one has to explain why a washing machine is not intelligent (relatively few humans can wash clothes) whereas a cat-recognizing program is (virtually every human, no matter how dumb, can recognize a cat, and so can countless animals). Statistically, it would seem that washing clothes should be more special than recognizing that a cat is a cat.

The current excitement about machines is due to the fact that (it is claimed) they are beginning to perform tasks that were exclusive to human beings. This is actually a very weak claim: the first washing machine was capable of performing a task that had been exclusive to human beings until the day before. Implicit in these claims is the idea that there is something that makes some tasks qualitatively more "special" than washing clothes, but it is difficult to articulate what this "special" quality would be. What is truly unique/special about human intelligence? Each machine performs for us a task that we used to do manually. Which tasks are so "special" that they deserve to be called "intelligent" is far from agreed upon.

And, finally, machines that resemble human beings (that smile, cry, walk and even say a few words) have existed for a long time and they are usually sold in toy stores and aimed at children. We can certainly create more sophisticated toys, like toys that recognize cats, but the claim that these toys will have anything to do with human intelligence needs some explaining.

Mind Uploading and Digital Immortality

Of all the life extension technologies proposed so far, none has captured the imagination of the Machine Intelligence crowd more than mind uploading. Somehow the connection between the Singularity and digital immortality was made: at some point those super-intelligent machines must be able to perform one great task for us, upload our entire self and "become" us. Couple it with the immortality of the "cloud" (see later), and your "self" becomes

immortal, being downloaded and uploaded from one release of the Singularity to the next one for the rest of time.

This is another idea that gave rise to a quasi-religious cult/movement, "transhumanism". The original prophet was probably Fereidoun "FM-2030" Esfandiary who wrote "Are You a Transhuman?" (1989) and predicted that "in 2030 we will be ageless and everyone will have an excellent chance to live forever". He died from pancreatic cancer (but was promptly placed in cryonic suspension).

In fact, the first major magazine for transhumanism predated his book: the first issue of Extropy magazine was published by Max More and Tom Morrow in 1988.

The technology of uploading a human mind into a computer was first explored by a geneticist, George Martin, in "A Brief Proposal on Immortality" (1971). He foresaw that some day computers would be so powerful that they will be able to do everything that a brain can do. Therefore why not simply port our brains to computers and let the computers do the job. Needless to say, philosophers are still arguing whether that "mind" would still be "me" once uploaded into software instead of gray matter.

That vision became more realistic in the 1990s with the explosion of the World-wide Web. Then paleontologist Gregory Paul in collaboration with mathematician Earl Cox spoke about cyber-evolution that would create non-human minds in "Beyond Humanity" (1996), including the idea of immortal brain carriers to replace our mortal bodies. In the days when television was still influential William Gibson, the science-fiction writer who had invented the term "cyberspace" ("Burning Chrome", 1982), contributed to popularize the concept by scripting an X-Files episodes, "Kill Switch" (1998), in which a man uploads his mind into cyberspace.

Then came the deluge with books such as Richard Doyle's "Wetwares - Experiments in PostVital Living" (2003) exploring all sorts of technologies of immortality. Every year the vision of what Martine Rothblatt calls "mindclones", implemented in "mindware"

(the software for consciousness), has to be updated to the latest computer platform.

In 2012 a Russian tycoon, Dmitry Itskov, pretty much summarized the vision of the immortality field: firstly, build brain-machine interfaces so that a human brain can control a robotic body; secondly, surgically transplant the human brain into a robotic body; and, finally, find a method to achieve the same result without the gory surgical operation, i.e. a way to upload a person's mind into the robotic body or, for that matter, into just about anything.

This program is predicated on the assumption that "i" am entirely in my brain, and that my body is simply a vehicle for my "i" to survive, but such body might as well be replaced by some other material substrate. The brain is disposable, according to this view: the brain is merely the organ of the body designated to host the processes that construct the "i", but the "i" truly is only those processes, which, luckily for us, turn out to be information-based processes, which, luckily for us, can be easily transplanted into the kind of information-processing machines that we started building in the 1940s and that are getting more and more powerful, rapidly approaching the capacity required to simulate the entirety of those brain processes.

This movement has revived the project of whole brain emulation. Ray Kurzweil and others now believe that artificial general intelligence will be achieved first via whole brain emulation. The basic idea is to construct a complete detailed software model of a human brain so that the hardware connected to that software will behave exactly like the human would (which includes answering the question "is it really you?" with a "yes").

But first one needs to map the brain, which is not trivial. In 1986 John White's and Sydney Brenner's team mapped the brain of the millimeter-long worm *Caenorhabditis Elegans* (302 neurons and 7000 synapses). As far as i know, that is still the only brain that we have fully mapped. And it took twelve years to complete that relatively simple "connectome". The term was coined in Olaf Sporns' "The Human Connectome, a Structural Description of the

Human Brain" (2005). A connectome is the map of all the neural connections in a brain. In 2009, a few years after the success of the Human Genome Project, the USA launched the Human Connectome Project to map the human brain. The problem, however, is not on the same scale. The entire human genome is represented by about a few gigabytes of data. Cellular biologist Jeff Lichtman and Narayanan Kasthuri estimated that a full human connectome would require one trillion gigabytes ("Neurocartography", 2010). Furthermore, we all share (roughly) the same genome, whereas each brain is different. The slightest mistake and, oops, they may upload the brain of someone else instead of yours.

Once we are able to map brains, we will need to interface those brains with machines. This may actually come sooner. In 1969 the Spanish neurophysiologist Jose Delgado implanted devices in the brain of a monkey and then sent signals in response to the brain's activity, thus creating the first bidirectional brain-machine-brain interface. In 2002 John Chapin debuted his "roborats", rats whose brains were fed electrical signals via a remote computer to guide their movements. His pupil Miguel Nicolelis achieved the feat of making a monkey's brain control a robot's arm. In 2008 the team made the monkey control a remote robot (in fact, located in another continent).

By the time science is capable of uploading your mind to cyberspace most of us will probably be dead, and with us our brains. That disturbing thought predated the very science we are talking about. Robert Ettinger's book "The Prospect of Immortality" (1962) is considered the manifesto of "cryonics", the discipline of preserving brains by freezing them. It was actually cryonics that started the "life extension" movement. In 1964 another founding father, Evan Cooper, launched the Life Extension Society (LES). In 1972 Fred Chamberlain, a space scientist at the Jet Propulsion Laboratory, founded the Alcor Society for Solid State Hypothermia (ALCOR), now called Alcor Life Extension Foundation, to enter that business.

The similarities with the most successful organized religions are too obvious to be overlooked. The end of the world is coming in the form of the Singularity, but, not to worry, we will all be resurrected in the form of mind uploads made possible by the super-machines of that very Singularity. The only difference with the old religions is that people from previous ages are dead for good, forever, as we don't have their brains to upload anymore. But then maybe those super-human machines will find a way to resurrect the dead too.

Machine Immortality and the Cloud

The other implicit assumption in the scenario of mind uploading is that these superhuman machines, capable of self-repairing and of self-replicating, will live forever.

That would represent a welcome change from what we are used to. The longest life expectancy for an electrical machine probably belongs to refrigerators, that can last longer than a human generation. Most appliances die within a decade. Computers are the most fragile of all machines: their life expectancy is just a few years. Their "memories" last less than human memory: if you stored your data on a floppy disc twenty years ago, there is probably no way for you to retrieve them today. Your CDs and DVDs will die before you. And even if your files survived longer, good luck finding an application that can still read them. Laptops, notepads and smartphones age increasingly faster. The life expectancy of machines seems to be decreasing ever more rapidly. And that's not to mention that they are alive only for as long as they are plugged into an electrical outlet (battery life can be as little as a few hours); and that they are much more vulnerable than humans to "viruses" and "bugs".

One has to be an inveterate optimist to infer from the state of the art that increasingly mortal and highly vulnerable computer technology is soon to become immortal.

Of course, the Singularity crowd will point out the "cloud", where someone else will take care of transferring your data from one dying

storage to a newer one and translating them from one extinct format to a newer one. Hopefully some day the cloud will achieve, if not immortality, at least the reliability and long lifespan of our public libraries, where books have lasted millennia.

Having little faith in software engineers, i am a bit terrified at the idea that some day the "cloud" will contain all the knowledge of the human race: one "bug" and the entire human civilization will be wiped out in one second. It is already impressive how many people lose pictures, phone numbers and e-mail lists because of this or that failure of a device. All it takes is that you forget to click on some esoteric command called "safely remove" or "eject" and an entire external disc may become corrupted. If and when super-intelligent machines come, i fear that they will come with their own deadly viruses, just like human intelligence came (alas) with the likes of flue pandemics, AIDS (Acquired ImmunoDeficiency Syndrome) and SARS (Severe Acute Respiratory Syndrome). And that's not to mention the likelihood of intentional cyber terrorism (i'm not sure who's getting better at it: the cryptographers who are protecting our data or the hackers who are stealing them) and of "malware" in general. If today they can affect millions of computers in a few seconds, imagine what the risk would be the day that all the knowledge of the world is held in the same place, reachable in nanoseconds. The old computer viruses were created for fun by amateurs. We are entering the age in which "cyber crime" will be the domain of super-specialists hired by corporations and governments. A computer virus was designed to be visible: that was the reward for its creator. The cyber crime that is coming soon will be invisible... until it's too late. Remember when only fire could destroy your handwritten notes on paper? And it was so easy to make photocopies (or even manual copies) of those handwritten notes? We found the Dead Sea Scrolls two thousand years after they had been written (on a combination of vellum and papyrus), and the Rosetta stone is still readable after 2,200 years. I wonder how many data that we are writing today can still be found two thousand years from now in the "cloud".

There is hope, however. I suspect that the widely publicized cyber-attacks of 2013, presumably sponsored by a large Asian state, will have the effect to wake up the cloud visionaries. Dozens of universities, corporations and government agencies were affected in the USA (other countries may still have to find out that they were cyber-attacked). The cyber-spies managed to crack passwords and steal thousands of documents. Hackers will keep getting more and more sophisticated and, when armed with powerful computers provided by rich governments, will be able to enter any computer that is online and access its contents (and possibly destroy them). In the old days the only way for a spy to steal a document was to infiltrate a building, search it, find the safe where the documents were being held, crack open the safe or bribe someone, duplicate the documents, flee. This was dangerous and time consuming. It could take years. Today a hacker can steal thousands if not millions of documents while comfortably sitting at her desk, and in a fraction of a second. The very nature of digital files makes it easy to search and find what you are looking for. I suspect that the future may see a return of paper (or some other physical format): the only way to make your files safe from hacking is to print them and then delete them from all computers. The hacker who wants to steal those files is now powerless, and has to be replaced by a traditional thief who has to break into your house, a much more complicated affair. Cyber-experts now admit that anything that you write in digital form and store on a device that is directly or indirectly connected to the Internet will, sooner or later, be stolen. Or destroyed. When, in march 2013, the websites of JPMorgan Chase and then American Express were taken offline for a few hours after being attacked by the Izz ad-Din al-Qassam Cyber Fighters, cyber-security expert Alan Paller warned that cyber-attacks are changing from espionage to destruction. A malware to destroy (digital) information on a large scale would be even easier to manufacture than the malware Stuxnet (unleashed in 2010 probably by Israel and the USA) that damaged about one thousand

centrifuges used to enrich nuclear material at Iran's nuclear facility in Natanz.

I also feel that "knowledge" cannot be completely abstracted from the medium, although i find it hard to explain what the difference is between knowledge stored in Socrates' mind, knowledge stored in a library and knowledge stored in a "cloud". A friend, co-founder of one of the main cloud-computing providers, told me: "Some day we'll burn all the books". Heinrich Heine's play "Almansor", written a century before Hitler, has a famous line: "Where they burn books, they will ultimately burn people too". Unlike most predictions about machine intelligence, that is one prediction that came true.

Corollary: Digital Media Immortality

If you want to turn yourself into data, instead of flesh and bones, and hope that this will make you immortal, you have a small technical problem to solve.

As you well know from your Christmas shopping, the capacity of computer storage media (for the same price) increases rapidly. That's the good news. The bad news is that its longevity has been decreasing, and significantly decreasing if you start from way back in time. The life expectancy of paper and ink is very long in appropriate conditions. The original storage media for computers, punched paper tapes and punch cards, are still readable 70 years later: unfortunately, the machines that can read them don't exist anymore, unless you have access to a computer museum. By comparison, the life expectancy of magnetic media is very, very short. Most people born before 1980 have never seen a magnetic tape except in old sci-fi movies. It was introduced with the first commercial computer, the Eckert-Mauchly's UNIVAC I, in 1951. Today most magnetic tapes store terabytes of data. They last about 20-30 years. Nobody knows how long the multiplatter disks from the mainframes of the 1960s lasted because they got out of fashion before we could test their lifespan. Floppy discs are magnetic disks, the most common type of which had a capacity of 1.44 megabytes

or 2 megabytes. The 8" floppy disks of the 1970s and the 5.25" floppy disks of the 1980s are given a life expectancy of 3-5 years by those who never used them, but those like me who still have them know that at least half of them are still working 30 years later. The "hard disks" that replaced them (and that today can easily hold a terabyte, i.e. a million times more data) may last longer, but they need to spin in order to be read or written, and spinning-disk hard drives don't last long: they are mechanical devices that are likely to break long before the magnetic layer itself deteriorates, especially if you carry them around (in other words, if you use them).

Music was stored on magnetic tapes, and later on cassettes, that would still work today if mass-market magnetic tape players still existed, although they would probably not sound too good, and on vinyl records, that definitely still play today if you didn't scratch them and used appropriate cartridges on your turntable like i did. My cassettes from the 1970s still play ok. Video was stored on VHS tapes, that still play today (i have about 300 of them), but, again, colors and audio may not look/sound so good after years of playing on a VCR (if you can still find a VCR).

Then came the optical generation. Rewritable optical discs are much less reliable for data storage than read-only optical discs that you buy/rent at music or video stores because they are physically made of different materials (the film layer degrades at a faster rate than the dye used in read-only discs). The jury is still out on optical media, but, as far as storing your data goes, the Optical Storage Technology Association (OSTA) estimates a lifespan of 10-25 years for compact discs (CDs), that typically held 650 megabytes (or the equivalent of 700 floppy disks), and digital video discs (DVDs), that typically held 4.7 gigabytes. However, in practice, optical devices are much more likely to get damaged because very few people store their discs in optimal conditions. Just leaving them on a desk unprotected may greatly shorten their lives just like anything else that you look at (optical is optical).

Now we live in the age of solid-state media, devices that don't have moving parts and that can store several gigabytes on a very

small device, like USB flash drives ("thumb" drives) and secure-digital cards ("flash cards"). They are generally less (not more) reliable than hard drives, and the manufacturers themselves don't expect them to last longer than about eight years.

And that's not to mention the quality of the recording: digital media are digital, not analog. You may not tell the difference because your ears are not as good as the ears of many animals, but the digital music on your smartphone is not as accurate a recording as the vinyl record of your parents or the 78 RPM record of your grandparents. A digital recording loses information. The advantage, in theory, is that the medium is less likely to deteriorate as you use it: magnetic tape degrades every time it passes by the magnetic head of a cassette player or a VCR, and the grooves of LPs do not improve when the cartridge of the turntable rides on them. The advantage of the old media, however, is that they "degraded": they didn't simply stop working. Digital files are either perfect or don't work, period. My old VHS tapes lost some of the color and audio fidelity, but i can still watch the movie. Many of my newer DVDs stop in the middle of the movie, and there is no way to continue. (I am also greatly annoyed by the difficult of rewinding/forwarding a DVD or pinpointing a frame of the movie, something that can be easily done on a VHS tape: this is possibly the first "regress" in history for random access, a feature introduced by the Romans when they switched from the scroll to the codex).

On the other hand, microfilms are estimated to last 500 years: that is a technology that was introduced by John Benjamin Dancer in 1839, and first used on a large scale in 1927 by the Library of Congress of the USA (that microfilmed millions of pages in that year).

You can tell that the plot remains the same: larger and larger storage, but perhaps less and less reliable.

Note that all of this is very approximate: search for the longevity of free neutrons, and you'll readily find it (14'42"), but if you search for a scientific answer to the question of storage media longevity, you will not find it. That's how advanced the science of storage media is.

And, even if your media could potentially last a long time, when is the last time you saw a new computer model with a 5 1/4" Floppy Disk? Even optical drives (CD, DVD) are disappearing as I type these words, and your favorite flash memory may become obsolete before this book goes out of print. And even if you still find a machine with a drive for your old media, good luck finding the operating system that has a file system capable of reading them. And even if you find both the hard drive and the operating system that can read them, good luck finding a copy of the software application that can read the data on them (e.g., GEM was the popular slide presentation software in the heydays of floppy discs). This is a field in which "accelerating progress" (in physical media, operating systems and viewing applications) has been consistently hurting data longevity, not extending it.

Ethical Intermezzo: The Moral Consequences of Human Immortality

One reason behind creating machine immortality is to hope that we can upload ourselves into the immortal machine and therefore become immortal ourselves.

Whether this is technically feasible or not, the idea that immortality can be achieved in this life has non-trivial consequences. If you believe that immortality is granted in the afterlife, you will do everything you can to obtain it in the afterlife (which typically means obeying the instructions that a god gave humans to achieve abovesaid immortality); but if you believe that immortality is granted in this life, you will do everything you can to obtain it in this life. A person who believes that immortality will be granted in the afterlife based on her good deeds will promptly sacrifice her life to save someone else or to fight a dangerous disease in Africa or to provide her children with a better future; but a person who believes that immortality is granted in this life has literally no motivation to risk her life to save someone else's life, nor any motivation to risk her life in Africa nor (ultimately) any

motivation to care for her own children. Once you are dead, you are dead, therefore the only meaning that your life can have is not to die. Ever. Under any circumstances. No matter what. Stay alive. Immortality in this life is the only thing that matters. That will be the new morality of a society that believes in immortality here now.

Philosophical Intermezzo: Do We Really Want Intelligence at All?

Intelligence is messy. When we interact with human beings, we have to consider their state of mind besides our immediate goal. We may only need a simple favor, but it makes a huge difference whether our interlocutor is happy or sad, on vacation or asleep, has just lost a close relative or been injured in an accident, angry at us, busy with many practical chores, etc. Whether our interlocutor is capable or not of performing that favor for us may be a secondary factor compared with whether she is in the mental condition of doing it and doing it right now. On the other hand, when we deal with a dumb machine, the only issue is whether the machine is capable of performing the task or not. If it is, and the power chord is plugged into the power outlet, it will. It won't complain, it won't ask us for a cigarette, it won't spend ten minutes gossiping about the neighbors, it won't comment on the government or the soccer game.

It may seem a paradox, but, as long as machines are dumb, they are easy and painless to interact with. They simply do what we ask them to do. No whims. No complaints. No formalities.

The complication that comes with intelligent beings is that they are subject to moods, feelings, opinions, intentions, motives, etc. There is a complicated cognitive apparatus at work that determines the unpredictable reaction of an intelligent being when you ask even the simplest of questions. If your friend is mad at you, even the simplest question "What time is it?" might not get an answer. On the other hand, if you use the right manners at the right time, a complete stranger may do something truly important for you. In

many cases it is crucial to know how to motivate people. In some cases that is not enough (if the person is in a bad mood for reasons that are totally independent of your will). Human beings are a mess. Dealing with them is a major project. And that's not to mention the fact that human beings need to sleep, get sick, go on vacation, and even take lunch breaks. In Western Europe they are often on strike.

Compare with dumb machines that have no feelings, no motives, no intentions, no opinions. They simply do what you ask. For example, the automatic teller machine hands you money at any time of the day or night any day of the year. Wherever the intelligent being has been replaced by a dumb machine, the interaction is simpler. (We realize that this is not the case whenever intelligence is required and then we get frustrated at being unable to talk with a human being, but this is happening less and less frequently as more and more exceptions are removed from highly structured environments).

The reason that automated customer support has replaced human beings in so many fields is not only that it is cheaper to operate by the provider but also that it is preferred in the majority of cases by the majority of customers. The honest truth is that very few of us enjoy waiting for an operator to tell us "Hello? How are you? Isn't it a beautiful day? How can i help you?" Most of us prefer to press digits on a phone keypad. The truth is that most customers are happy if we remove the complication of dealing with human beings.

When i worked in the corporate world, my top two frustrations were the secretary and the middle management. Dealing with the secretary (especially in unionized Italy) required superior psychological skills: say the wrong word in the wrong tone, and she'll boycott your entire day. Most middle managers were mediocre and mostly slowed down things, seemingly paid mainly to kill great ideas. The only way to get important things done quickly was, again, to use the art of psychology: befriend them, chat with them, find out what motivated them, offer them rides home, hang

out with them. Don't i wish that my colleagues and my secretary had been heartless robots?

And, let's face it, we often don't have the patience for human interactions that involves protocols of behavior. We are often happy when good manners are replaced by cold mechanic interactions shaped by goals and restrained by laws. Hence we do not really want machines with human intelligence, i.e. we don't want them to have emotions, to be verbose, to deceive, to plead, etc. One goal of inventing machines is precisely to remove all of that, to remove that inefficient, annoying, time-consuming quality of "humanity".

We removed the human/intelligent element from many facets of ordinary life because the truth is that in most cases we don't want to deal with intelligent beings. We want to deal with very dumb machines that will perform a very simple action when we press a button.

I'll let psychologists and anthropologists study why this trend towards less and less humane interactions, but the point here is that intelligence comes at a high price: intelligence comes with feelings, opinions, habits, and a lot of other baggage. You can't have intelligence without that baggage.

Analog vs Digital

Most machines in the history of human civilization were and are analog machines, from the waterwheel to your car's engine. A lot of the marvel about computers comes from the fact that they are digital devices: once you digitize texts, sounds, images, films and so forth the digital machine can perform at incredible speed operations that used to take armies of human workers or specialists armed with expensive specialized machines. Basically, digitizing something means reducing it to numbers, and therefore the mind-boggling speed of computers in performing calculation gets automatically transferred to other fields, such as managing texts, audio and video. The "editing" feature is, in fact, one of the great revolutions that came with the digital world. Previously, it was

difficult and time-consuming to edit anything (text, audio, photos, video, etc). Filing, editing and transmitting are operations that have been dramatically revolutionized by technological progress in digital machines and by the parallel process of digitizing everything, one process fueling the other.

Now that television broadcasts, rented movies, songs and books are produced and distributed in digital formats, i wonder if people of the future will even know what "analog" means. Analog is any physical property whose measurable values vary in a continuous range. Everything in nature is analog: the weight of boulders, the distance between cities, the color of cherries, etc. (At microscopic levels nature is not so analog, hence Quantum Theory, but that's another story). Digital is a physical property whose measurable values are only a few, for example zero and one. I don't know any digital device that is not binary. The digital devices of today can typically handle only two values: zero and one. Hence, de facto, "digital" and "binary" mean the same thing in 2013. Numbers other than zero and one can be represented by sequences of zeroes and ones (e.g. a computer internally turns 5 into 101). Texts, sounds and images are represented according to specific codes (such as ASCII, MP3 and MP4) that turn texts, sounds and images into strings of zeroes and ones.

The easiest way to visualize the difference between analog and digital is to think of the century-old bell-tower clocks (with the two hands crawling between the 12 Roman numerals) and the digital clock (that simply displays the time as hour/minutes).

When we turn a property from analog to digital we enable computers to deal with it. Therefore you can now edit, copy and email a song (with simple commands) because it has been reduced to a music file (to a string of zeroes and ones). In the old days the only way for an ordinary person to "copy" a song was to tape it with a tape recorder or equivalent. However, audiophiles still argue whether digital "sounds" the same as analog. I personally think that it does (at today's bit rates) but the stubborn audiophile has a point: whenever we digitize an item, something is lost. The digital clock

that displays "12:45" does not possess the information of how many seconds are missing to 12:46. Yesterday's analog clock contained that information in the exact position of the minute hand. That piece of information may have been useless (and obtainable only by someone equipped with a magnifying glass and a pocket calculator) but nonetheless the device had it. The music file is not the song: when the musicians performed it, they were producing an analog object. Once that analog object is turned into a digital file, an infinite number of details have been lost. The human ear is limited and therefore won't notice (except abovesaid stubborn audiophiles). We don't mind because our senses can only experience a limited range of sound and visual frequencies. And we don't mind because amazing features become available with digital files, for example the ability to improve the colors of a photograph.

When machines carry out human activities, they are "digitizing" those activities; and they are digitizing the "mental" processes that lie behind those activities. In fact, machines can manage those human activities only after humans digitized (turned into computer files) everything that those human activities require, for example maps of the territory.

Using digital electronic computers to mimic the brain is particularly tempting because it was discovered that neurons work like on/off switches. They "fire" when the cumulative signal that they receive from other neurons exceeds a certain threshold value, otherwise they don't. Binary logic seems to lie at the very foundation of human thinking. In fact, as early as 1943, Warren McCulloch, in cooperation with Walter Pitts, described mathematically an "artificial" neuron that can only be in one of two possible states. A population of artificial neurons can then be connected in a very intricate network to mimic the way the brain works. When signals are sent into the network, they spread to its neurons according to the simple rule that any neuron receiving enough positive signals from other neurons sends a signal to other neurons. It gets better: McCulloch and Pitts proved that a network of binary neurons is fully equivalent to a Universal Turing Machine.

There is, however, a catch: McCulloch's binary neurons integrate their input signals at discrete intervals of time, rather than continuously as your brain's neurons do. Every computer has a central clock that sets the pace for its logic, but the brain relies on asynchronous signaling because there is no synchronizing central clock. If you get into more details of how the brain works, there are more "analog" processes at work, and there are analog processes inside the neuron itself (which is not just a switch). One could argue that the brain is regulated by the body's internal clocks (that regulate every function, from your heart to your vision) and therefore the brain behaves like a digital machine. This is true, but it's like saying that nothing in nature is truly analog because, at a very fine level of detail (at the microscopic level, if necessary) everything is made of discrete objects all the way down to quarks and leptons. In practice, a brain uses a lot more than zeroes and ones; a computer can only deal with zeroes and ones. As tempting as it is to see the brain as a machine based on binary logic, the difference between the human brain and any computer system (no matter how complex the latter can become) is that a computer is way more digital than a brain. We know so little about the brain that it is difficult to estimate how many of its processes involve a lot more than on/off switching (and can be considered "analog" for practical purposes), but a safe guess is that we're talking several hundreds. Despite the illusion created by the McCulloch-Pitts neuron, for practical purposes a computer is a digital machine and a brain is not.

There might be a reason if a brain operates at 10-100 Hz whereas today's common microprocessors need to operate at 2-3 Gigahertz (billions of Hz), hundreds of millions of times faster to do a lot less; and if human brains consume about 20 Watts and can do a lot more things than a supercomputer that consume millions of Watts. Biological brains need to be low-power consumption machines or they would not survive. There are obviously principles at work in a brain that have eluded computer scientists.

Carver Mead's "neuromorphic" approach to machine intelligence is not feasible for the simple reason that we don't know how the brain works. Based upon the Human Genome Project (that successfully decoded the human genome in 2003), the USA launched the "Brain Initiative" in april 2013 to map every neuron and every synapse in the brain.

There are also government-funded projects to build an electronic model of the brain: Europe's Human Brain Project and the USA's Systems of Neuromorphic Adaptive Plastic Scalable Electronics (SYNAPSE), sponsored by the same agency, DARPA, that originally sponsored the Arpanet/Internet. Both Karlheinz Meier in Germany and Giacomo Indiveri in Switzerland are toying with analog machines. The signaling from one node to the others better mimics the "action potentials" that triggers the work of neurons in the human brain and requires much less power than the ones employed in digital computers. SYNAPSE (2008) spawned two projects in California, one run by Narayan Srinivasa at Hughes Research Laboratories (HRL) and the other one run by Dharmendra Modha at IBM's Almaden Labs in Silicon Valley. The latter announced in 2012 that a supercomputer was able to simulate 100 trillion synapses from a monkey brain, and in 2013 unveiled its "neuromorphic" chip True North (not built according to the traditional John von Neumann architecture).

Teaser: Machine Ethics

If we ever create a machine that is a fully-functioning brain totally equivalent to a human brain, will it be ethical to experiment on it? Will it be ethical to program it?

Summarizing and The Timeframe for Artificial General Intelligence

In April 2013 i saw a presentation at Stanford's Artificial Intelligence Lab by the team of Kenneth Salisbury in collaboration with the Willow Garage about a robot that can take the elevator and walk upstairs to buy a cup of coffee. This implies operations that are trivial for humans: recognizing that a transparent glass door is a door (not just a hole in the wall and never mind the reflection of the robot itself in the glass), identifying the right type of door (rotating, sliding or automatic), finding the handle to open the door, realizing that it's a spring-loaded door so it doesn't open as easily as regular doors, finding the elevator door, pressing the button to call the elevator, entering the elevator, finding the buttons for the floors inside an elevator whose walls are reflective glass (therefore the robot keeps seeing reflections of itself), pressing the button to go upstairs, locating the counter where to place the order, paying, picking up the coffee, and all the time dealing with humans (people coming out of the door, sharing the space in the elevator, waiting in line) and avoiding unpredictable obstacles; if instructions are posted, read the instructions, understand what they mean (e.g. this elevator is out of order or the coffee shop is closed) and change plan accordingly. The robot got it right. It took the robot 40 minutes to return with the cup of coffee. It is not impossible. It is certainly coming. I'll let the experts estimate how many years it will take to have a robot that can go and buy a cup of coffee upstairs in all circumstances (not just those programmed by the engineer) and do it in 5 minutes like humans do. The fundamental question, however, is whether this robot can be considered an intelligent being because it can go and buy a cup of coffee or it is simply another kind of appliance (we don't normally assume that a dishwasher is as intelligent as a human being just because it is capable of washing dishes).

It will take time (probably much longer than the optimists claim) but some kind of "artificial intelligence" is indeed coming. How soon

depends on your definition of artificial intelligence. One of the last things that John McCarthy wrote before dying was: "We cannot yet characterize in general what kinds of computational procedures we want to call intelligent" (2007). To start with, it is wise to make a distinction between artificial intelligence and A.G.I. (Artificial General Intelligence). Artificial intelligence is coming very soon if you don't make a big deal of it: we are just using a quasi-religious term for "automation", a process that started with the waterwheels of ancient Greece if not earlier. Search engines (using very old fashioned algorithms and a huge number of very modern computers housed in "server farms") will find an answer to any question you may have. Robots (thanks to progress in manufacturing and to rapidly declining prices) will become pervasive in all fields, and become household items, just like washing machines and toilets; and eventually some robots will become multifunctional (just like today's smartphones combine the functions of yesterday's watches, cameras, phones, etc; and, even before smartphones, cars acquired a radio and an air conditioning unit, and planes acquired all sorts of sophisticated instruments). Millions of jobs will be created to take care of the infrastructure required to build robots, and to build robots that build robots, and ditto for search engines, websites and whatever comes next. Some robots will come sooner, some will take centuries. And miniaturization will make them smaller and smaller, cheaper and cheaper. At some point we will be surrounded for real by Neil Stephenson's "intelligent dust" (see his novel "Diamond Age"), i.e. by countless tiny robots each performing some function that used to be exclusive to humans. If you want to call it "artificial intelligence", suit yourself.

However, "artificial general intelligence" should be more than the ability to solve some specific "intelligent" task. Most of the systems that are being touted as A.I. today (the ones based on neural networks that solve one specific problem thanks to massive number crunching) are not what i would call A.I., but, if that's what you call A.I., then your AGI depends on which tasks you want it to be capable of. Someone has to list whether AGI requires being able to

ride a bus in Zambia and to exchange money in Haiti or whether it only requires the ability to sort out huge amounts of data at lightning speed or what else. Once we have that list, we can ask the world's specialists in each of those fields to make reasonable predictions. Making predictions about AGI without having a clear definition of what constitutes an AGI is as scientific as making predictions about the coming of Jesus.

This is an old debate. Many decades ago the founders of computational mathematics (Alan Turing, Claude Shannon, Norbert Wiener, John von Neumann and so forth) discussed which tasks can become "mechanic", i.e. performed by a computing machine, i.e. what can and what cannot be computed, i.e. what can be outsourced to a machine and what kind of machine it has to be. As far as I know, today's computers that perform today's neural-network algorithms are still Universal Turing Machines, subject to the theorems proven for those classes of machines. If the complex architectures of parallel computers are no longer Universal Turing Machines, then someone should prove similar theorems for these new classes of machines. If they are UTMs, then Alan Turing's original work still rules. The whole point of inventing (conceptually) the Turing Machine in 1936 was to prove whether a general algorithm to solve the "halting problem" for all possible program-input pairs exists, and the answer was a resounding "no" (there is always at least one program that will never terminate). And in 1951 Henry Gordon Rice generalized this conclusion with an even more formidable statement, "Rice's Theorem" (that any nontrivial property about the language recognized by a Turing machine is undecidable). In other words, it is theoretically proven that there is a limit to what machines can "understand", no matter how much progress there will be, if they are Universal Turing Machines (as all of today's digital devices are).

Nonetheless, by employing thousands of these machines, the "brute force" approach has achieved sensational feats such as machines that can beat chess champions and recognize cats. I call this "number-crunching" kind of A.I. because its success depends

on computational mathematics. If, like me, you believe that the human mind does not work that way, and therefore that A.I. should not be about "number crunching", then we are still waiting for a theoretical breakthrough. I suspect that the current success in the "number-crunching" kind of A.I. is slowing down (not accelerating) research in higher-level architectures (in what to me is the real meaning of "human intelligence"). If i can fix a car without knowing anything about cars, why bother to figure out how cars work? The success in (occasionally) recognizing cats, beating chess champions and the like is indirectly reducing the motivation to understand how the human mind (or, for that matter, the chimp's mind, or even a worm's mind) manages to recognize so many things in split seconds. The success in building robots that perform this or that task with amazing dexterity is indirectly reducing the motivation to understand how the human mind can control the human body in such sophisticated manners.

Simply telling me that Artificial Intelligence and robotic research will keep producing better and smarter devices (that are fundamentally not "intelligent" the way humans are) does not tell me much about the chances of a breakthrough towards a different kind of machine that will match (general) human intelligence.

I don't know what a breakthrough looks like, but i know what it doesn't look like. The machine that beat the world champion of chess was programmed with knowledge of virtually every major chess game ever played and was allowed to run millions of logical steps before making any move. That obviously put the human contender at a huge disadvantage. Even the greatest chess champion with the best memory can only remember so many games. The human player relies on intuition and creativity, whereas the machine relies on massive doses of knowledge and processing. Shrink the knowledge base that the machine is using to the knowledge base that i have and limit the number of logical steps it can perform to the number of logical steps that the human mind can perform before it is timed out, and then we'll test how often it wins against ordinary chess players, let alone world champions.

Having a computer (or, better, a huge knowledge base) play chess against a human being is like having a gorilla fight a boxing match with me: i'm not sure what conclusion you could draw from the result of the boxing match about our respective degrees of intelligence.

I wrote that little progress has been made in Natural Language Processing. The key word is "natural". Machines can actually speak quite well in unnatural language, a language that is grammatically correct but from which all creativity has been removed: "subject verb object - subject verb object - subject verb object - etc." The catch is that humans don't do that. If i ask you ten times to describe a scene, you will use different words each time.

Language is an art. That is the problem. How many machines do we have that can create art? How far are we from having a computer that switches itself on in the middle of the night and writes a poem or draws a picture just because the inspiration came? Human minds are unpredictable. And not only adult human minds: pets often surprise us, and children surprise us all the time. When is the last time that a machine surprised you? (other than surprise you because they are still so dumb). Machines simply do their job, over and over again, with absolutely no imagination. We haven't been able to instill any degree of creativity in machines.

Here is what would constitute a real breakthrough: a machine that has only a limited knowledge of all the chess games ever played and is allowed to run only so many logical steps before making a move and can still consistently beat the world champion of chess. That machine will have to use intuition and creativity. That's a machine that would probably wake up in the middle of the night and write a poem. That's a machine that would probably learn a human language in a few months just like even the dumbest children do. That is a machine that would not translate "'Thou' is an ancient English word" into "'Tu' e` un'antica parola Inglese", and that will not stop at a red traffic light if it creates a dangerous situation.

I suspect that this will require some major redesigning of the very architecture of today's computers. For example, a breakthrough

could be a transition from digital architectures to analog architectures. Another breakthrough could be a transition from silicon (never used by Nature to construct intelligent beings) to carbon (the stuff of which all natural brains are made). And another one, of course, could be the creation of a machine that is self-conscious.

Today it is commonplace to argue that in the 1960s A.I. scientists gave up too quickly on neural networks. My gut feeling is that in the 2000s we gave up a bit too quickly on the symbolic processing program. Basically, we did to the knowledge-based approach what we had done before to the connectionist approach: in the 1970s neural networks fell into oblivion because knowledge-based systems were delivering practical results... only to find out that knowledge-based systems were very limited and that neural networks were capable of doing more. My guess is that there was nothing wrong with the knowledge-based approach: it is just that we never figured out an adequate way to represent human knowledge. Representation is one of the oldest problems in philosophy, and I don't think we got any closer to solving it now that we have powerful computers. The speed of the computer does little to fix a wrong theory of representation. Then we decided that the knowledge-based approach was wrong and we opted for neural networks (deep learning and the likes). And neural networks have proven very good at simulating specialized tasks: each neural network does one thing well, but doesn't do what every human, even the dumbest one, and even animals, do well: use the exact same brain to carry out millions (potentially an infinite number) of different tasks.

As it stands, predictions about the future of (really) intelligent machines are predictions about something that doesn't exist. You can ask a rocket scientist for a prediction for when a human being will travel to Pluto: that technology exists and one can speculate what it will take to use that technology for that specific mission. On the contrary, my sense is that, using current technology, there is no way that we can create a machine that is even remotely capable of

our routine cognitive tasks. The technology that is required does not yet exist, and will not exist until the day of that breakthrough.

It is difficult to predict the future because we tend to predict one future instead of predicting all possible futures. Nobody (as far as I know) predicted that the idea of expert systems would become irrelevant in most fields because millions of volunteers would post knowledge for free on something called World-wide Web accessible by anybody equipped with something called "browser". That was one possible future but there were so many possible futures that nobody predicted this one. By the same token, it is hard to predict what will make sense in ten years, let alone in 50 years.

What if 3D printing and some other technology makes it possible for ordinary people to create cheap gadgets that solve all sorts of problems. Why would we still need robots? What if synthetic biology starts creating alternative forms of life capable of all sorts of amazing functions. Why would we still need machine intelligence? There is one obvious future, the one based on what is around today, in which machines would continue to multiply and evolve. There are many other futures in which digital electronic machines would become irrelevant because of something that does not exist today.

And if you think that A.G.I. requires consciousness, read the debate on whether that is possible at all: it may never happen.

Anders Sandberg and Nick Bostrom, authors of "Whole Brain Emulation" (2008), conducted a "Machine Intelligence Survey" (2011) that starts with a definition of what an Artificial Intelligence should be: a system "that can substitute for humans in virtually all cognitive tasks, including those requiring scientific creativity, common sense, and social skills." My estimate for the advent of such a being is roughly 200,000 years: the timescale of natural evolution to produce a new species that will be at least as intelligent as us. If Artificial Intelligence has to be achieved by incremental engineering steps starting from the machines that we have today, my estimate about when a machine will be able to carry out a conversation like this one with you is: "Never". I am simply

projecting the progress that i have witnessed in Artificial Intelligence (very little and very slow) and therefore i obtain an infinite time required for humans to invent such a machine. And then, again, we'd probably have a lengthy discussion about what "all cognitive tasks" really means. For example, leaving out consciousness from the category of cognitive tasks is like leaving out Beethoven from the category of musicians simply because we can't explain his talent.

The technology of batteries was born in 1859 when Gaston Plante invented the first rechargeable battery (a few years later Georges Laclanche invented the forerunner of today's alkaline battery). Since then the performance of batteries doubled every 70 years: quite a slow pace. Then in 1991 Japan's Sony released the first commercial lithium-ion battery and performance quickly quadrupled. Even lithium-ion batteries might not truly represent a breakthrough because we are still making and using batteries the same way that Plante and Laclanche used them, but that was real progress (in understanding how to store and release energy), not simply simulation by "brute force" of real progress.

Don't ask me what the breakthrough will be in A.I. If i knew it, i wouldn't be wasting my time writing articles like this one. But i have a hunch it has to do with recursive mechanisms for endlessly remodeling internal states: not data storage, but real "memory".

For historians a more interesting question is what conditions may foster a breakthrough. In my opinion, it is not the abundance of a resource (such as computing power or information) that triggers a major paradigm shift but the scarcity of a resource. For example, James Watt invented the modern steam engine when and because Britain was in the middle of a fuel crisis (caused by the utter deforestation of the country). For example, Edwin Drake discovered petroleum ("oil") in Pennsylvania when and because whale oil for lamps was becoming scarce. Both innovations caused an economic and social revolution (a kind of "exponential progress") that completely changed the face of the world. The steam engine created an economic boom, reshaped the landscape, revolutionized

transportation, and dramatically improved living conditions. Oil went on to provide much more than lighting to the point that the contemporary world is (alas) addicted to it. I doubt either revolution would have happened in a world with infinite amounts of wood and infinite amounts of whale oil. The very fact that computational power is becoming an infinite inexpensive resource makes me doubt that it will lead to a breakthrough in Artificial Intelligence. Water power was widely available to Romans and Chinese, and they had the scientific know-how to create machines propelled by water; but the industrial revolution had to wait more than one thousand years. One reason (not the only one but a key one) why the Romans and the Chinese never started an industrial revolution is simple: they had plentiful cheap labor (the Romans had slaves, the Chinese emperors specialized in mobilizing masses). Abundance of a resource is the greatest deterrent to finding an alternative to that resource. If "necessity is the mother of ingenuity", as Plato said, then abundance is the killer of ingenuity. We live in the age of plentiful computational power. What to some looks like evidence that super-human machine intelligence is around the corner looks to me like evidence that our age doesn't even have the motivation to try.

On the other hand, i have seen astonishing (quasi exponential) progress in Biotechnology, and therefore my estimate for when Biotech will be able to create an "artificial intelligence" is very different: it could already happen in one year. And my estimate of when Biotech might create a "superhuman" intelligence is also more optimistic: it could happen in a decade. I am simply basing my estimates on the progress that i have witnessed over the last 50 years; which might be misleading (again, most technologies eventually reach a plateau), but at least this one has truly been "accelerating" progress. In the countries where stem cell research is legal a new science is already thriving, regenerative medicine, the science of producing tissues and organs that are absolutely normal (albeit "artificial" in that they were grown in a laboratory) such as Madeline Lancaster's "organoids" ("Cerebral Organoids Model

Human Brain Development And Microcephaly”, 2013). It would be interesting to discuss how Biotech might achieve this feat: will it be a new being created in the laboratory, or the intended or the accidental evolution of a being, or a cell by cell replica of the human body? But that's for another book.

The real deal is the digital to biological conversion that will increasingly allow biologists to create forms of life. That is indeed a “breakthrough”. Machines will remain a tool (that every generation will brand “intelligent” and every generation will expect to get more “intelligent”) but one of their applications, the biotech application, is likely to have the biggest impact on the future of life on this planet.

If by “artificial intelligence” you simply mean a machine that can do something (not everything) that we can do (like recognizing cats or playing chess), but not “everything” that we can do (both the mouse and the chess player do a lot of other things), then all machines and certainly all appliances qualify. Some of them (radio, telephone, television) are even forms of superhuman intelligence because they can do things that human brains cannot do.

Definitions do matter: there is no single answer to the questions “when will an artificial general intelligence be built” and “when will superhuman intelligence happen”. It depends on what we mean by those words. My answer can be “it's already here” or “never”.

As i wrote, machines are making us somewhat dumber (or, better the environments we design for automation make us dumber), and there is an increasing number of fields (from arithmetic to navigation) in which machines are now “smarter” than humans not only because machines got smarter but also because humans have lost skills that they used to have. If i project this trend to the future, there is a serious chance that humans will get so much dumber that the bar for artificial intelligence will be lower and therefore artificial intelligence more feasible than it is today; and “superhuman” intelligence may then happen, but it should really be called “subhuman” intelligence.

Nick Bostrom wrote that the reason A.I. scientists have failed so badly in predicting the future of their own field is that the technical

difficulties have been greater than they expected. I don't think so. I think those scientists had a good understanding of what they were trying to build. The reason why "the expected arrival date [of artificial intelligence] has been receding at a rate of one year per year" (Bostrom's estimate) is that we keep changing the definition. There never was a proper definition of what we mean by "artificial intelligence" and there still isn't. No wonder that the original A.I. scientists were not concerned with safety or ethical concerns: of course, the machines that they had in mind were chess players and theorem provers. That's what "artificial intelligence" originally meant. Being poor philosophers and poor historians, they did not realize that they belonged to the centuries-old history of automation, leading to greater and greater automata. And they couldn't foresee that within a few decades all these automata would become millions of times faster, billions of times cheaper, and would be massively interconnected. The real progress has not been in A.I. but in miniaturization. Miniaturization has made it possible to use thousands of tiny cheap processors and to connect them massively. The resulting "intelligence" is still rather poor, but its consequences are much more intimidating.

The statistical method that has become popular in Artificial Intelligence during the 2000s is simply an admission that previous methods were not wrong but simply difficult to apply to all problems in general. This new method, like its predecessors, can potentially be applied to every kind of problem until scientists will admit that... it cannot. The knowledge-based method proven inadequate for recognizing things and was eventually abandoned (nothing wrong with it at the theoretical level). The traditional neural networks proved inadequate for just about everything because of their high computational costs. In both cases dozens of scientists had to tweak the method to make it work in a narrow and very specific problem domain. When generalized, the statistical methods in vogue in the 2010s turn out to be old-fashioned mathematics such as statistical classification and optimization algorithms. These might indeed be more universal than previous A.I. methods but, alas,

hopelessly computational-resource intensive. It would be ironic if Thomas Bayes' theorem of 1761 turned out to be the most important breakthrough in Artificial Intelligence. Unfortunately, it is easy to find real-world problems in which the repeated application of that theorem leads to a combinatorial explosion of the space of potential solutions that is computationally intractable. We are now waiting for the equivalent of John Hopfield's "annealing" algorithm that in 1982 made neural networks easier to implement. That will make this Bayesian kind of Artificial Intelligence go for a little longer, but i am skeptical that this will lead to a general A.I. The most successful algorithms used in the 2010s to perform machine translation require virtually no linguistic knowledge. The very programmer who creates and improves the system has no knowledge of the two languages being translated into each other: it is only a statistical game. I doubt that this is how human interpreters translate one language into another, and i doubt that this approach will ever be able to match human-made translations, let alone surpass it.

Donald Knuth's famous sentence that A.I. seems better at emulating "thinking" than at emulating the things we do without thinking is still true; and it contains a larger truth. The real hard problem is that we don't know how we do the vast majority of things that we do, otherwise philosophers and psychologists would not have a job. A conversation is the typical example. We do it effortlessly. We shape strategies, we construct sentences, we understand the other party's strategy and sentences, we get passionate, we get angry, we try different strategies, we throw in jokes and we quote others. Anybody can do this without any training or education. And now, by comparison, check what kind of conversation can be carried out by the most powerful computer ever built. It turns out that most of the things that we do by "thinking" (such as proving theorems and playing chess) can be emulated with a simple algorithm (especially if the environment around us has been shaped by society to be highly structured and to allow only for a very small set of moves). The things that we do

without thinking cannot be emulated with a simple algorithm, if nothing else because even we don't know how we do them. We can't even explain how children learn in the first place. Artificial Intelligence scientists have a long way to go because they have a poor philosophical understanding of what humans do and a poor neurophysiological understanding of how humans do it.

As you may have guessed, someone like me who is skeptical about the present and the near future of Artificial Intelligence is not likely to predict that the Singularity is coming any time soon. But it is actually worse: i suspect that the Singularity may just be impossible. I cannot prove it mathematically, but this is what i feel intuitively. There is no generally accepted definition of intelligence, so i might as well introduce my own. Let's define intelligence as a property of a category of systems: the intelligence of a category is the set of all the systems that any system in that category could potentially build. For the purposes of this discussion, we will only focus on two types of categories: biological species and machines. In particular, human intelligence (the intelligence of the category of humans) is the set of all the things that humans can potentially build. The great physicist Stephen Hawking may not be able to build anything because he is confined to a wheelchair but, potentially, he too could build cars and skyscrapers. In particular, machine intelligence in 2013 is zero because there is no machine that can build (of its own will) anything. This definition makes sense to me because we don't rank animals based on survival skills anymore: we rank them based on the skills that allow a species to rule over other species; and that's largely the ability to build things. Better sniffing skills do not create superior intelligence, graphene (a new material invented a decade ago) does. "Superior" is the species that is capable of building more and better systems. The Singularity would be a generation of machines that is superior to humans in building other machines.

My thesis is that a system cannot create a more intelligent system. All the systems that humans can create are less intelligent than humans. In particular, human intelligence is superior or equal to the intelligence of a Universal Turing Machine, and therefore of

all machines that are Turing machines such as computers. If a Turing machine, say T1, builds another kind of machine, say T2, either T2 is another Turing machine like T1 or it is less intelligent than T1. And so on: T3 created by T2 will be equally or less intelligent than T2.

One could object that species evolved, and now we have more intelligent species "created" by less intelligent ones. But it is not species that create new species: it is natural selection. If a force like natural selection can apply to machines, then, sure, machines can "evolve into" (not "build") more intelligent machines.

I cannot prove mathematically that the Singularity is impossible, but maybe a thermodynamical explanation might explain where this intuition is coming from. The "building skills" can be expressed in terms of thermodynamical systems. Every living being and every machine is a thermodynamical system. A human being is a thermodynamical system. I absorb energy from the environment (oxygen, water, food, solar energy) and turn it into both energy that propels my organs and excrements that I discard. This thermodynamical system can build other thermodynamical systems that we call machines, but the laws of Thermodynamics limit what these products can be, and, in particular, according to my definition of "intelligence", they cannot be more intelligent than me. Anybody who thinks otherwise is basically introducing a mystical religious element, that we are more than a thermodynamical system.

Provocations aside, the superhuman machine will better arrive within 10,000 years otherwise, even if we save the planet from the current suicidal trends, it will be too late: in 1983 physicist Brandon Carter introduced the "Doomsday Argument" later popularized by philosopher John Leslie in his book "The End of the World: The Science and Ethics of Human Extinction" (1996). This was a simple mathematical theorem based on Bayesian Probability Theory demonstrating that we can be 95% certain that we are among the last 95% of all the humans ever to be born. It has been tweaked up and down by various dissenters but, unlike Singularity Science, it sits on solid mathematical foundations: in fact, the exact same

foundations (Bayesian reasoning) of today's most popular Deep Learning neural-net systems.

Intermezzo: We may Overestimate Human Brains

Is "general intelligence" possible at all? We think of ourselves as a "general intelligence", capable of anything, but are we? Is our brain truly a general-purpose machine? Can our brain's circuits truly adapt to any task? The evidence from neuroscience is still scant, but it could be that the brain is simply a society of brains, each specialized in a specific form of intelligence, and the set of all those "intelligences" is the most we can do. The visual cortex doesn't seem to have much in common with the auditory cortex, and even less in common with the regions of the brain that account for language. It could well be that we are not so "general" as we'd like to be. A common criticism of Artificial Intelligence is that it creates machines that are "intelligent" in only one field, whereas our (natural) intelligence is "fluid": we can reuse something learned to solve a task to solve a completely different kind of task. The verdict is still out whether this is true or not. I am not sure that when I learn to better calibrate with eyesight, I can transfer that learning to, say, orient myself better in the wilderness.

The Real Future of Computing

In 1988 Mark Weiser proposed the vision of a future in which computers will be integrated into everyday objects ("ubiquitous computing") and these objects will be connected with each other. This became known as the "Internet of Things" after 1998 when two MIT experts in Radio-frequency identification (RFID), David Brock and Sanjay Sarma, figured out a way to track products through the supply chain with a "tag" linking to an online database.

The technology is already here: sensors and actuators have become cheap enough that embedding them into ordinary objects will not significantly increase the price of the object. Secondly, there

are enough wireless ways to pick up and broadcast data that it is just a matter of agreeing on some standards like Near Field Communication (NFC). Monitoring these data will represent the next big wave in software applications. Facebook capitalized on the desire of people to keep track of their friends: people own many more "things" than friends, spend more time with things than with friends, and do a lot more with things than with friends.

At the same time, one has to be aware that the proliferation of digital control also means that we will live in an increasingly surveilled world because there will be a record of everything that has happened and that is happening. Machines indirectly become spies. In fact, your computer (desktop, laptop, notepad, smartphone) is already a sophisticated and highly accurate spy that records every move you make: what you read, what you buy, whom you talk to, where you travel, etc. All this information is on your hard disk and can easily be retrieved by forensic methods.

The focus of computer science is already shifting towards collecting, channeling, analyzing and reacting to billions of data arriving from all directions. Luckily, the "Internet of Things" will be driven by highly structured data.

The data explosion is proceeding faster than Moore's law, i.e. than the increase in processing speed: exploring data is becoming increasingly difficult with traditional John von Neumann computer architectures (that are designed for calculations). IBM's Watson employs multiple statistical engines to learn features about the data, and is in fact (in 2013) being applied to the medical field (cancer research). Watson does not make decisions, only data analysis for scientific discovery, especially useful to discover unexpected correlations among data that would probably elude human beings trained not to see those. It is what used to be called "a decision support tool", an example of augmented intelligence, helping humans acquire more knowledge. Just like with other decision support tools, human analysts could perform the same task, but not at the same speed.

If I am skeptic about the creation of an agent that will be a general artificial intelligence, I am very aware that we are rapidly creating a sort of global intelligence as we connect more and more software and this giant network produces all sorts of positive feedback loops. This network is already out of control and every year gets harder to control.

The Future of Miniaturization

If I am right and the widely advertised progress in machine intelligence is mainly due to rapid progress in miniaturization and cost reduction, then it would be more interesting to focus on the future of miniaturization. Whatever miniaturization achieves next is likely to determine the "intelligence" of future machines.

While IBM's Watson was stealing the limelight with its ability to answer trivia questions, others at IBM were achieving impressive results in Nanotechnology. In 1989 Don Eigler's team at IBM's Almaden Research Center, using the scanning tunneling microscope built in 1981 by Gerd Binnig and Heinrich Rohrer, carried out a spectacular manipulation of atoms that resulted in the atoms forming the three letters "IBM". In 2012 Andreas Heinrich's team (in the same research center) stored one magnetic bit of data in 12 atoms of iron, and a byte of data in 96 atoms; and in 2013 that laboratory "released" a movie titled "A Boy and His Atom" made by moving individual atoms. This ability to trap, move and position individual atoms using temperature, pressure and energy could potentially create a whole new genealogy of machines.

The Moral Issue: Who's Responsible for a Machine's Action?

During the 2000s drones and robotic warfare stepped out of science-fiction novels/movies into reality. These weapons represent the ultimate example of how machines can be used for sophisticated albeit dubious purposes. If I accidentally kill three

children, i will feel guilty for the rest of my life and perhaps commit suicide. But who feels guilty if the three children are killed by mistake by a drone that was programmed 5000 kms away by a team using Google maps, Pakistani information and Artificial Intelligence software, a strike authorized by the president of the USA in person? The beauty of delegating tasks to machines is that we decouple the action from the perpetrator, or, at least, we dilute the responsibility so much that it becomes easier to "pull the trigger" than not to pull it. And what if the mistake was due to malfunctioning software? Will the software engineer feel guilty? She may not even learn that there was a bug in her piece of software, and, if she does, she may never realize that the bug caused the death of three children.

Of course, this process of divorcing the killing from the killer started at least in World War I with the first aerial bombings (a practice later immortalized by Pablo Picasso, when it still sounded horrible, in his painting "Guernica") and that happened precisely because humans were using machines (the airplanes) to drop the bombs on invisible citizens instead of throwing grenades or shooting guns against visible enemies. The killer will never know nor see the people he killed.

What applies to warfare applies to everything else. The use of machines to carry out whatever action basically relieves the "designer" of responsibility.

The same concept can be applied, for example, to surgery: if the operation performed by a machine fails and the patient dies, who is to blame? The team that controlled the machine? The company that built the machine? The doctor who prescribed the use of that specific machine? I suspect that none of these will feel particularly guilty. There will simply be another machine that will mechanically add one to a statistical number of failed procedures. "Oops: you are dead". That will be the reaction of society to a terrible incident.

You don't need to think of armed drones to visualize the problem. Think of a fast-food chain. You order at a counter, then you move down the counter to pay at the cash register, and then you hang out

by the pick-up area. Eventually some other kid will bring you the food that you ordered. If what you get is not what you ordered, it is natural to complain with the kid who delivered it; but he does not feel guilty (correctly so) and his main concern is to continue his job of serving the other customers who are waiting for their food. In theory, you could go back to the ordering counter, but that would imply either standing in line again or upsetting the people who are in line. You could summon the manager, who was not even present when the incident happened, and blame him for the lousy service. The manager would certainly apologize (it is his job), but even the manager would be unable to pinpoint who is responsible for the mistake (the kid who took the order? the chef? the pen that wasn't writing properly?)

In fact, many businesses and government agencies neatly separate you from the chain of responsibility so that you will not be able to have an argument with a specific person. When something goes wrong and you get upset, each person will reply "i just did my job". You can blame the system in its totality, but in most cases nobody within that system is guilty or gets punished. And, still, you feel that the system let you down, that you are the victim of an unfair treatment.

This manner of decoupling the service from the servers has become so pervasive that younger generations take it for granted that often you won't get what you ordered.

You can already see how ordinary people deflect responsibility for their failures. For example, people who are late for an appointment routinely blame their gadgets. For example, "The navigator sent me to the wrong address" or "The online maps are confusing" or "My cell phone's batteries died". In all of these cases the implicit assumption is that you are not responsible, the machine is. The fact that you decided to use a navigator (instead of asking local people for directions) or that you decided to use those online maps (instead of the official government maps) or that you forgot to recharge your cell phone doesn't seem to matter anymore. It is taken for granted

that your life depends on machines that are supposed to do the job for you and, if they don't, it is not your fault.

There are many other ethical issues that are not obvious. Being a writer who is bombarded with copyright issues all the time, here is one favorite. Let us imagine a future in which someone can create an exact replica of any person. The replica is just a machine, although it looks and feels and behaves exactly like the original person. You are a pretty girl and a man is obsessed with you. That man goes online and purchases a replica of you. The replica is delivered by mail. He opens the package, enters an activation code and the replica starts behaving pretty much like you would. Nonetheless, the replica is, technically speaking, just a toy. It has no feelings/emotions, it just simulates the behavior that your feelings/emotions would cause. Then this man proceeds to have sex with that replica of you against its will. This is a toy bought from a store, so it is perfectly legal to do anything the buyer wants to do with it, even to rape it and even to kill it. I think you get the point: we have laws that protect this very sentence that you are reading from being plagiarized and my statements from being distorted, but no law protects a full replica of us.

Back to our robots capable of critical missions: since they are becoming easier and cheaper, they are likely to be used more and more often to carry out these mission-critical tasks. Easy, cheap and efficient: no moral doubts, no falling asleep, no double crossing. The temptation to use machines instead of humans in more and more fields will be too strong to resist.

I wonder if it is technology that drives the process (which would already be bad in itself, since the process is supposed to drive technology and not viceversa) or is it the desire to be relieved of moral responsibility that drives the adoption of a technology and then the process that justifies that technology.

One wonders whether society is aiming for the technology that minimizes our responsibility instead of aiming for the technology that maximizes our effectiveness, let alone aiming for the technology that maximizes our accountability.

The Dangers of Machine Intelligence: Machine Credibility

One of my hiking buddies works for a major mapping service. Whenever we find a "bug" in the system (not unusual when you are out of the main traffic routes), he immediately emails headquarters to fix the problem. On the contrary, ordinary users of that service (the most popular service on today's smartphones) not only insist on using it even when there is someone in the car who knows the route very well (and they will stop the car if the navigation system stops working), but they also tend to defend the service even when faced with overwhelming evidence that it took us to the wrong place or via a ridiculous route.

In September 2013 i posted on Facebook that YouTube was returning an ad about (sic) pooping girls when i looked for "Gandhi videos". An incredible number of people wrote back that the ad was based on my search history. But i was not logged into YouTube, Gmail or any other product. A friend (who has been in the software industry all his life) then wrote "It doesn't matter, Google knows". It was pointless to try and explain that if you are not logged in, the software (whether Google, Bing or anything else) does not know who is doing the search (it could be a guest of mine using my computer, or it could be someone who just moved into my house using the same IP address that i used to have). And it was pointless to swear that i had never searched for pooping girls! (For the last week or so i had been doing a research to compile a timeline of modern India). Anyway, the point is not that i was innocent, but that an incredible number of people were adamant that Google knows that i am the one doing that search (only Google or all software?) It reminded me of the Catholic priest in elementary school: "God knows!" Maybe we're going down the same path. People will believe that software can perform miracles when in fact most software has bugs that make it incredibly stupid. This might be very similar to what happened with the birth of religions. (Next they

started burning at the stakes the heretics like me who kept repeating "God is just a software bug!")

The faith that an ordinary user places in a digital gadget wildly exceeds the faith that its very creators place in it. If i make a mistake just once giving directions, i lose credibility for a long time; if the navigation system makes a mistake, most users will simply assume it was an occasional glitch and will keep trusting it. The tolerance for mistakes seems to be a lot higher when it comes to machines.

People tend to believe machines more than they believe humans, and, surprisingly, seem to trust machine-mediated opinions better than first-hand opinions from an expert. For example, they will trust the opinions expressed on social media like Yelp more than they will trust the opinion of the world's expert on restaurants. They believe their navigation system more than they believe someone who has spent her entire life in the neighborhood. Therefore they are likely to believe strongly a machine that has a huge knowledge base.

Imagine a machine that broadcasts false news, for example that an epidemic is spreading around New York killing people at every corner. No matter what the most reputable reporters write, people will start fleeing New York. Panic would rapidly spread, from city to city, amplified by the very behavior of the millions of panicking citizens (and, presumably, by all the other machines that analyze, process and broadcast the data fed by that one machine).

Generally speaking, we are a lot less smart than we think, and we can easily be fooled by humans. When we use a computer, we seem to become even more vulnerable. Think of how successful "spam" is, or even of how successful the ads posted by your favorite search engine and social media are. If we were smarter, those search engines and social media would rapidly go out of business. They thrive because millions of people click on those links.

The more "intelligent" software becomes, the more likely that people trust it. Unfortunately, at the same time the more "intelligent"

it becomes, the more capable of harming people it will be. It doesn't have to be "intentionally" evil: it can just be a software bug, one of the many that software engineers routinely leave behind as they roll out new software releases that most of us never asked for.

Drone strikes seem to enjoy the tacit support of the majority of citizens in the USA. That tacit support arises not only from military calculations (that a drone strike reduces the need to deploy foot soldiers in dangerous places) but also from the belief that drone strikes are accurate and will mainly kill terrorists. However, drone strikes that the USA routinely hails as having killed terrorists are often reported by local media and eyewitnesses in Pakistan, Afghanistan, Yemen and so on as having killed a lot of harmless civilians, including children. People who believe that machines are intelligent are more likely to support drone strikes. Those who believe that machines are still very dumb are very unlikely to support drone strikes. The latter (including me) believe that the odds of killing innocents are colossal because machines are so dumb and are likely to make awful mistakes (just like the odds that the next release of your favorite operating system has a bug are almost 100%). If everybody were fully aware of how inaccurate these machines are, i doubt that drone programs would exist for much longer.

In other words, i am not so much afraid of machine intelligence as of human gullibility.

The Dangers of Machine Intelligence: Speed Limits for Machines?

In our immediate future i don't see the danger that future machines will be conceptually difficult to understand (superhuman intelligence), but i do see the danger that future machines will be so fast that controlling them will be a major project in itself. Again, it may all be about semantics: computational speed does matter. We

already cannot control a machine that computes millions of times faster than our brain, and this speed will keep increasing in the foreseeable future. That's not to say that we cannot understand what the machine does: we perfectly understand the algorithm that is being computed. In fact, we wrote it and fed it into the machine. It is computed at a much higher speed than the smartest mathematician could. When that algorithm leads to some automatic action (say, buying stocks on the stock market), the human being is left out of the loop and has to accept the result. When thousands of these algorithms (each perfectly understandable by humans) are run at incredible speed by thousands of machines interacting with each other, humans have to trust the computation. It's the speed that creates the "superhuman" intelligence: not an intelligence that we cannot understand, but an intelligence vastly "inferior" to ours that computes very quickly (and that's all it does). The danger is that nobody can make sure that the algorithm was designed correctly, especially when it interacts with a multitude of algorithms.

The only thing that could be so fast is another algorithm. I suspect that this problem will be solved by introducing the equivalent of speed limits: algorithms will be allowed to compute at only a certain speed, and only the "cops" (the algorithms that stop algorithms from causing problems) will be allowed to run faster.

The Dangers of Machine Intelligence: Criminalizing Common Sense

There is also something disturbing about the machines that intelligent humans are building with the specific mandate to overcome the desires of intelligent humans. Stupid machines in charge of making sure that human intelligence does not interfere with rules and regulations are becoming widespread in every aspect of life. I'll take a simple example because i find it even more telling than the ones that control lives at higher and more sinister levels. I live in the San Francisco Bay Area, one of the most technologically advanced regions in the world. In one of the most prestigious

universities of the Bay Area we hold evening events. Because of a world-famous fog, the weather is chilly (if not cold) in the summer, especially in the evening. Nonetheless, the computers have been programmed to provide air conditioning throughout campus for as long as there are people at work, no matter how cold it is getting outside. People literally bring sweaters and even winter coats at these evening classes. Never mind the total waste of energy; the interesting point is that nobody knows anymore how to tell the machines to stop doing that (insert "smiley" here). After months of trying different offices, we still are "not sure who else to contact about it" (quoting the head of a department in the School of Science) "apparently it is very difficult to reset the building's thermostat". This is the real danger of having machines run the world. I don't think any of us would call a thermostat "intelligent" when it directs very cold air into a room during a very cold evening. In fact, we view it as... beyond stupid. However, it is very difficult for the wildly more intelligent race that created it to control its behavior according to common sense. The reason is that this incredibly stupid machine was created to overcome the common sense with which more intelligent beings are equipped. Think about it and probably your computer-controlled car, some of your computer-controlled appliances and computer-controlled systems around you often prohibit you from performing actions that common sense and ordinary intelligence would demand, even when they work perfectly well, and, in fact, precisely because they work perfectly well.

I am afraid of the millions of machines that will populate human society with the specific goal of making sure that humans don't use common sense but simply follow the rules.

The Dangers of Machine Intelligence: You Are a Budget

Another danger is that what will truly increase exponentially is the current trend to use computing power as a sales tool. The reason that people are willing to accept the terms and conditions of e-

commerce websites is that these companies have been very good at concealing what they do with the information that they collect about you. As a former Facebook research scientist (Jeff Hammerbacher) put it: "The best minds of my generation are thinking about how to make people click ads". It's not only the best minds but also the best machines. Artificial Intelligence techniques are already used to gather information on you (what used to be called "spy on you") for the purpose of targeting you with more effective sales strategies. The more information, the smarter the machine has to be to analyze it; and the smarter the machine, the more it can find out about you. The original purpose of the World-wide Web was not to create a world in which smarter software controls every move you make online and uses it to tailor your online experience; but that is precisely what it risks becoming. Computer science is becoming the discipline of turning your life into somebody else's business opportunity.

Why We Shouldn't Talk About the Singularity

The first and immediate reason why obsessive discussions about the coming of machine super-intelligence and human immortality are harmful is that they completely miss the point.

We live in an age of declining innovation. Fewer and fewer people have the means or the will to become the next Edison or Einstein. The great success stories in Silicon Valley (Google, Facebook, Apple) are of companies started by individuals with very limited visions that introduced small improvements over existing technologies. Entire nations (China and India, to name the obvious ones) are focusing on copying, not inventing.

Scholars from all sorts of disciplines are discussing the stagnation of innovation. A short recent bibliography: Tyler Cowen's e-book "The Great Stagnation" (2010) by an economist; Neal Stephenson's article "Innovation Starvation" (2011) by a sci-fi writer; Peter Thiel's article "The End of the Future" (2011) by a Silicon Valley venture capitalist; Max Marmor's "Reversing The Decline In Big Ideas"

(2012) by another Silicon Valley entrepreneur; Jason Pontin's "Why We Can't Solve Big Problems" (2012) by a technology magazine editor; Rick Searle's article "How Science and Technology Slammed into a Wall and What We Should Do About It" (2013) by a political scientist.

(If you didn't even know that there is a debate about the stagnation of innovation, you just proved my point).

Then there is the fundamental issue of priorities. The hypothetical world of the Singularity distracts us from the real world. The irrational exuberance about the coming Singularity distracts a lot of people from realizing the dangers of unsustainable growth, dangers that may actually wipe out all forms of intelligence from this planet.

Let's assume for a second that climate scientists like Paul Ehrlich and Chris Field (to name two I met in person at Stanford) are right about the coming apocalypse. Their science is ultimately based on the same science that happens to be right about what that bomb would do to Hiroshima (as unlikely as Einstein's formula may look), that is right about what happens when you speak in that rectangular device (as unlikely as it may seem that someone far away will hear your voice), that is right about what happens when someone broadcasts a signal in that frequency range to a box sitting in your living room (as unlikely as it may seem that the box will then display the image of someone located far away), that is right about what happens when you turn on that switch (as unlikely as it is that turning on a switch will light up a room); and it's the same science that got it right on the polio vaccine (as unlikely as it may look that invisible organisms cause diseases) and many other incredible affairs.

The claims about the Singularity, on the other hand, rely on a science (Artificial Intelligence) whose main achievement has been to win a chess game. One would expect that whoever believes wholeheartedly in the coming of the Singularity would believe tenfold stronger that the human race is in peril.

Let's assume for a second that the same science that has been right on just about everything it predicted is also right on the

consequences of rapid climate change and therefore the situation is exactly the opposite of the optimistic one based mostly on speculation depicted by A.I. science: the human race may actually go extinct before it even produces a single decent artificial intelligence.

In about one century the Earth's mean surface temperature has increased by about 0.8 degrees. Since it is increasing faster today than it was back then, the next 0.8 degrees will come even faster, and there is widespread agreement that 2 degrees above what we have today will be a significant tipping point. Recall that a simple heat wave in summer 2003 led to 15,000 deaths in France alone. Noah Diffenbaugh and Filippo Giorgi (authors of "Heat Stress Intensification in the Mediterranean Climate Change Hotspot ", 2007) have created simulations of what will happen to the Earth with a mean temperature 3.8 degrees above today's temperature: it would be unrecognizable. That temperature, as things stand, is coming for sure, and coming quickly, whereas super-intelligence is just a theoretical hypothesis and, in my humble opinion, is not coming any time soon.

Climate scientists fear that we may be rapidly approaching a "collapse" of civilization as we know it. There is not one but several environmental crises. Some are well known: extinction of species (with unpredictable biological consequences, such as that declining populations of bees may pose a threat to fruit farms), pollution of air and water, epidemics, and, of course, anthropogenic (man-made) climate change. See the "Red List of Threatened Species" published periodically by the International Union for Conservation of Nature (IUCN). See the University of North Carolina's study "Global Premature Mortality Due To Anthropogenic Outdoor Air Pollution and the Contribution of Past Climate Change" (2013), that estimated air pollution causes the deaths of over two million people annually. A Cornell University study led by David Pimentel, "Ecology of Increasing Diseases" (2007), estimated that water, air and soil pollutions account for 40% of worldwide deaths. A 2004 study by the Population Resource Center found that 2.2 million

children die each year from diarrhea caused by contaminated water and food. And, lest we think that epidemics are a thing of the past, it is worth reminding ourselves that AIDS (according to the World Health Organization) has killed about 35 million people between 1981 and 2012, and in 2012 about 34 million people were infected with HIV (Human Immunodeficiency Virus, the cause of AIDS), which makes it the fourth worst epidemics of all times. Cholera, tuberculosis and malaria are still killing millions every year; and "new" viruses routinely pop up in the most unexpected places (Ebola, West Nile virus, Hantavirus, Avian flue, etc).

Some environmental crises are less advertised but no less terrifying. For example, global toxification: we filled the planet with toxic substances, and now the odds that some of them interact/combine in some deadly runaway chemical experiment never tried before are increasing exponentially every year. Many scientists point out the various ways in which humans are hurting our ecosystem, but few single out the fact that some of these ways may combine and become something that is more lethal than the sum of its parts. There is a "non-linear" aspect to what we are doing to the planet that makes it impossible to predict the consequences.

The next addition of one billion people to the population of the planet will have a much bigger impact on the planet than the previous one billion. The reason is that human civilizations have already used up all the cheap, rich and ubiquitous resources. Naturally enough, humans started with the cheap, rich and ubiquitous ones, whether forests or oil wells. A huge amount of resources is still left, but those will be much more difficult to harness. For example, oil wells have to be much deeper than they used to. Therefore one liter of gasoline today does not equal one liter of gasoline a century from now: a century from now they will have to do a lot more work to get that liter of gasoline. It is not only that some resources are being depleted, but even the resources that will be left are, by definition, those that are difficult to extract and use (a classic case of "diminishing margin of return").

The United Nations' "World Population Prospects" (2013) estimated that the current population of 7.2 billion will reach 9.6 billion by 2050, and population growth will mainly come from developing countries, particularly in Africa: the world's 49 least developed countries may double in size from around 900 million people in 2013 to 1.8 billion in 2050.

The bottom line of these arguments is that the collapse is not only coming, but the combination of different kinds of environmental problems makes it likely that it is coming even sooner than the pessimists predict and in a fashion that we cannot quite predict.

For the record, the environmentalists are joined by an increasingly diversified chorus of experts in all sorts of disciplines. For example, Jeremy Grantham who is an economist (managing 100 billion dollars of investments). His main point (see, for example, his 2013 interview on Charlie Rose's television program) is that the "accelerated progress" that the Singularity crowd likes to emphasize started 250 years ago with the exploitation of coal and then truly accelerated with the exploitation of oil. The availability of cheap and plentiful energy made it possible to defy, in a sense, the laws of Physics. Without fossil fuels the human race would not have experienced such dramatic progress in merely 250 years. Now the planet is rapidly reaching a point of saturation: there aren't enough resources for all these people. Keeping what we have now is a major project in itself, and those who hail the coming super-intelligence miss the point the way a worker about to get fired missed the point if he is planning to buy a bigger house.

We are rapidly running out of cheap resources, which means that the age of steadily falling natural resource costs is coming to an end. In fact, the price of natural resources declined for a century until about 2002 and then in just 5 or 6 years that price regained everything that it had lost in the previous century (i am still quoting Grantham). This means that we may return to the world of 250 years ago, before the advent of the coal (and later oil) economy, when political and economic collapses were the norm; a return to, literally, the ages of starvation.

It is not only oil that is a finite resource: phosphates are a finite resource too, and the world's agriculture depends on them.

Alternative energy and slowing population growth will have to be adopted worldwide and rapidly in order to avoid a collapse of civilization.

Population growth is indeed slowing down everywhere, but "overpopulation" is measured in terms of material resources: most developed countries are not overcrowded, not even crowded Singapore, because they are rich enough to provide a good life to their population; most underdeveloped countries are overcrowded because they can't sustain their population.

Another point is consumption, which is increasing precisely where population growth is declining: one billion Chinese who ride bicycles is not the same as one billion Chinese who drive cars, run A/C units and wrap everything in plastic. If you do it, why shouldn't they?

The very technologies that should improve people's lives (including your smartphone and the future robots that you are dreaming of) are likely to demand more energy, which for now comes mainly from the very fossil fuels that are leading towards a catastrophe and possibly to the demise of the human race.

All those digital devices will require more "rare earths", more coltan, more lithium and many other materials that are becoming scarcer.

We also live in the age of Fukushima, when the largest economies are planning to get rid of nuclear power, which is the only form of clean alternative energy as effective as fossil fuels. Does anyone really think that we can power all those coming millions of robots with wind turbines and solar panels?

Chris Field has a nice diagram (expanded in the 2012 special report of the Intergovernmental Panel on Climate Change titled "Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation") that shows "Disaster Risk" as a function of "Climate Change" and "Vulnerability" (shown, for example, at a seminar at the Energy Biosciences Institute in 2013). It is worth pondering the effects of robots, A.I. and the likes on that

equation. Climate change can be naturally occurring or anthropogenic (man-made). Two big factors account for vulnerability: development (that reduces vulnerability) and population (that increases it). Manufacturing machines (millions of machines) will have an impact on anthropogenic climate change; economic development comes at the cost of exploitation of finite resources; and, if high technology truly succeeds in increasing the longevity of the human race, the population will keep expanding. In conclusion, the race to create intelligent machines might exacerbate the risk of disasters before these super-intelligent machines can find a way to reduce it.

Economists such as Robin Hanson ("Economics Of The Singularity", 2008) have studied the effects of the agricultural, industrial and digital revolutions. Each caused an acceleration in economic productivity. The world's GDP may double every 15 years on average in this century. That's an impressive feat, but it's nothing compared with what would happen if machines could replace people in every single task. Productivity could then double even before we can measure it. The problem with that scenario is that the resources of the Earth are finite, and most wars have been caused by scarcity of resources. Natural resources are already strained by today's economic growth. Imagine if that growth increased ten fold, and, worse, if those machines were able to mine ten or 100 times faster than human miners. It could literally lead to the end of the Earth as a livable planet. In our age the biggest global problem is climate change, caused by humans. That is dangerous enough if it spins out of control. Now imagine a new global problem, caused by machines that rapidly multiply and improve, and basically use all of the Earth's resources within a few years.

Ehrlich calls it "growthmania": the belief that there can be exponential growth on a finite planet.

The counter-argument is that digital technology can be "cleaner" than the old technology. For example, the advent of email has dramatically reduced the amount of paper that is consumed, which has reduced the number of trees that we need to fell. It is also

reducing the amount of mail trucks that drive around cities to deliver letters and postcards. Unfortunately, in order to check email and text messages you need devices like laptops, notepads and smartphones. The demand for materials such as lithium and coltan has risen exponentially.

Technological progress in the internal combustion engine (i.e., in fuel-efficient vehicles), in hybrid cars, in electric cars and in public transportation is credited for the reduction in oil consumption since 2007. The truth is that 2008 marked the beginning of the Great Recession, which drastically reduced the number of commuter cars and commercial trucks on the road at any given time. New fuel economy standards enacted by the USA will reduce oil consumption by 2.2 million barrels per day by 2025, but at the same time oil consumption in China has increased by some 3328% between 1965 and 2006, and even that still represents only 8% of the world's oil consumption for the country that has 20% of the world's population. Asia Pacific as a whole has posted a 46% increase in oil consumption in the first decade of the 21st century. In 2000 oil consumption in China was 4.8 million bpd (barrels per day), or 1.4 barrels per person per year. In 2010 China's consumption had grown to 9.1 million bpd. China and India together have about 37% of the world's population. The rate of cars per person in China (0.09%) is almost 1/10th the one in the USA (0.8%) and in India is one of the lowest in the world (0.02%). Hence analysts such as Ken Koyama, chief economist at the Institute of Energy Economics Japan, predict that global petroleum demand will grow 15% over the next two decades ("Growing Oil Demand and SPR Development in Asia", 2013).

George Mitchell pioneered fracking in 1998, releasing huge amounts of natural gas that were previously thought inaccessible. Natural gas may soon replace oil in power stations, petrochemical factories, domestic heaters and perhaps motor vehicles. The fact that there might be plenty of this resource in the near future proves that technology can extend the life expectancy of natural resources,

but it does not change the fact that those resources are finite, and it might reduce the motivation to face the inevitable.

Technology is also creating a whole new biological ecosystem around us, a huge laboratory experiment never tried before. Humans have already experienced annihilation of populations by viruses. Interestingly, both the most famous ones took hold at a time of intense global trade: the plague of 1348 (the "black death") was probably brought to Europe by Italian traders who picked it up in Mongol-controlled regions at a time when travel between Europe and Asia was relatively common and safe; and the flu pandemic of 1918, that infected about 30% of the world's population and killed 50 million people, took hold thanks to the globalized world of the British and French empires and to World War I. The Ebola virus always makes news, but it kills its hosts so rapidly that it does not spread too far. On the other hand, the HIV is a lentivirus (slowly replicating retrovirus). It came out in the 1980s when the Western economies had become so entangled and, sure enough, it spread first to the USA, then to Europe, then eventually everywhere. By the end of 2012 AIDS had killed 35 million people worldwide.

We now live in the fourth experiment of that kind: the most globalized world of all times, in which many people travel to many places; and they do so very quickly. There is one kind of virus that could be worse than the previous ones: a coronavirus, whose genes are written in RNA instead of DNA. The most famous epidemics caused by a coronavirus was the Severe Acute Respiratory Syndrome (SARS): in February 2003 it traveled in the body of a passenger from Hong Kong to Toronto, and within a few weeks it had spread all over East Asia. Luckily both Canada and China were equipped to deal with it and all the governments involved did the right thing; but we may not be as lucky next time.

True: we don't have good models for any of these problems. One would hope that the high-tech industry invest as much into creating good computational models to save the human race as into creating ever more lucrative machines. Otherwise, before the technological singularity, we may enter an "ecological singularity".

Discussing super-human intelligence is a way to avoid discussing the environmental collapse that might lead to the disappearance of human intelligence. We may finally find the consensus to act on environmental problems only when the catastrophe starts happening. Meanwhile, the high-tech world will keep manufacturing, marketing and spreading the very items that make the problem worse (more vehicles, more electronic gadgets, and, soon, more robots); and my friends in Silicon Valley will keep boasting about the latest gadgets (the things that environmental scientists call "unnecessarily environmentally damaging technologies"), firmly believing that we are living in an era of accelerating progress.

Fans of high technology fill their blogs with news of ever more ingenious devices to help doctors, not realizing that the proliferation of such devices will require even more energy and cause even more pollution (of one sort or another). They might be planning a world in which we will have fantastic health care tools but we will all be dead.

I haven't seen a single roadmap that shows how technology will evolve in the next decades, leading up to the Singularity (to super-human intelligence). I have, instead, seen many roadmaps that show in detail what will happen to our planet under current trends.

Last but not least, we seem to have forgotten that a nuclear war (even if contained between two minor powers) would shorten the life expectancy of everybody on the planet, and possibly even make the planet uninhabitable. Last time I checked, the number of nuclear powers had increased, not decreased, and, thanks to rapid technological progress and to the electronic spread of knowledge, there are now many more entities capable of producing nuclear weapons.

The enthusiastic faith that Rome was the "eternal city" and the firm belief that Venice was the "most serene republic" did not keep those empires from collapsing. Unbridled optimism can be the most lethal weapon of mass destruction.

The Future of Human Intelligence

Maybe we should focus on what can make us (current Homo Sapiens people) more intelligent, instead of focusing on how to build more intelligent machines that will make our intelligence obsolete.

How you define "intelligence" makes a difference. Personally, I think it is not about memorizing or calculating, but about "creating". Creativity is what truly sets Homo Sapiens apart from other species.

There are two myths here that I never bought. The first one is that adults are smarter than children. Children perform an impressive feat in just a few years, acquiring an incredible amount of knowledge and learning an incredible portfolio of skills. They are also fantastically creative in the way they deal with objects and people. Teenagers are still capable of quick learning (for example, foreign languages) and can be very creative (often upsetting parents and society that expect a more orthodox behavior, i.e. compliance with rules). Adults, on the other hand, tend to live routine lives and follow whatever rules they are told to obey. The overwhelming evidence is that intelligence declines with age. We get dumber and less creative, not smarter and more creative; and, once we become dumb adults, we do our best to make sure that children too become as dumb as us.

Secondly, the people of the rich developed high-tech world implicitly assume that they are smarter than the people of the poor undeveloped low-tech world. In my opinion, nothing could be farther from the truth. The top of creativity is encountered in the slums and villages of the world. It is in the very poor neighborhoods that humans have to use their brain every single minute of their life to come up with creative and non-orthodox solutions, solutions that nobody taught them to problems that nobody studied before. People manage to run businesses in places where there is no infrastructure, where at any time something unpredictable can (and will) happen. They manage to sell food without a store. They manage to trade without transportation. When they obtain a tool, they often use it not for the purpose for which it was originally designed but for some other purpose. They devise ever new ways

to steal water, electricity, cable television and cellular phone service from public and private networks. They find ways to multiply and overlap the functions of the infrastructure (for example, a railway track also doubles as a farmer's market, and a police road-block becomes a snack stop). They help each other with informal safety networks that rival state bureaucracies (not in size or budget, but in effectiveness). The slums are veritable laboratories where almost every single individual (of a population of millions) is a living experiment (in finding new ways of surviving and prospering). There is no mercy for those who fail to "create" a new life for themselves every day: they stand no chance to "survive".

If one could "measure" creativity, I think the slums of any city would easily win over Silicon Valley.

These highly creative people yearn for jobs in the "white" economy, the economy of the elite that lives outside the slums. For that "white" economy they may perform trivial repetitive jobs (chauffeur, cashier, window washer); which means that they have to leave their creativity at home. The "white" economy has organized daily life in such a way (around "routines") that everybody is guaranteed to at least survive. The people of the slums use their brains only when they live and work in the slums. When they live or work outside the slums, they are required to stop being creative and merely follow procedures, procedures that were devised by vastly less creative people who would probably not survive one day in the slums. Routines maximize productivity precisely by reducing human creativity. Someone else has created, and the worker only has to perform, a series of predefined steps.

The routine, on the other hand, can "amplify" the effect of innovation. The innovation may be very small and very infrequent, but the effect of the routine performed by many workers (e.g., by many Silicon Valley engineers) is to make even the simplest innovation relevant for millions of individuals.

The creativity of slums and villages, on the other hand, is constant, but, lacking the infrastructure to turn it into routine, ends up solving only a small problem for a few individuals. The slums are

a colossal reservoir of creative energies that the world is wasting, and, in fact, suppressing.

In our age we are speeding up the process by which (rule-breaking) children become (rule-obeying) adults and, at the same time, we are striving to turn the creativity of the slums into the routine of factories and offices. It seems to me that these two processes are more likely to lead to a state of lower rather than higher intelligence for the human race.

I suspect that removing the unpredictable from life means removing the very essence of the human experience and the very enabler of human intelligence.

A Conclusion: Sociology Again

Humans have been expecting a supernatural event of some kind or another since prehistory. Millions of people are still convinced that Jesus will be coming back soon, and millions believe that the Mahdi will too. The Singularity risks becoming the new religion for the largely atheistic crowd of the high-tech world. Just like with Christianity and Islam, the eschatological issue/mission then becomes how to save oneself from damnation when the Singularity comes, balanced by the faith in some kind of resurrection. We've seen this movie before, haven't we?

P.S. What i did not Say

I want to emphasize what i did not write in this book.

I did not claim that an artificial general intelligence is impossible, (only that it requires a major revolution in the field); and i certainly did not argue that superhuman intelligence is not possible (in fact, i explained that it is already all around us); and i did not rail against technological progress (i simply lamented that its achievements and benefits are wildly exaggerated).

I did not write that technology makes you stupid. I wrote that rules and regulations make you stupid; and technology is used to

establish, enforce and multiply those rules and regulations (often giving machines an unfair advantage over humans).

I did not write that there has been no progress in Artificial Intelligence: there has been a lot of progress, but mainly because of cheaper, smaller and faster processors, and because of more structured environments in which it is easier to operate (both for humans and for machines).

I did not write that humans will never create an artificial intelligence. We have already created artificial intelligence programs that do many useful things, as well as some truly obnoxious ones (like displaying ads on everything you do online). The definition of "intelligence" is so vague that the very first computer (or the very first clock) can be considered an artificial intelligence. In fact, early computers were called "electronic brains", not "electronic objects". McCarthy was right when he complained that, whenever Artificial Intelligence achieves something, that "something" is no longer considered an example of artificial intelligence.

I did not write that Artificial Intelligence is useless. On the contrary, i think that it has helped neuroscience. Maybe the "enumeration" problem (the problem of enumerating all the intelligent tasks that are needed in order to achieve artificial general intelligence) is a clue that our own brain might not be "one" but a confederation of many brains, each specialized in one intelligent task.

Appendix: The Dangers of Clouding - Wikipedia as a Force for Evil

Since 2001, when Jimmy Wales started it, Wikipedia has grown to become the largest cooperative project in the history of the world, with 39 million articles in 250 languages (as of 2013). The jury, however, is still out on whether Wikipedia is a force for good or a force for evil.

The traditional debate over Wikipedia has focused on how much we can trust thousands of anonymous editors as opposed to the small team of highly decorated scholars who curate the traditional encyclopedia. Since scholars and erudite people in general are less likely to get into a brawl, the fear was that in the long run the "mob" would win, the frequent outcome in many facets of popular culture. That was pretty much the only concern when Wikipedia was just that: a substitute for the encyclopedia.

However, the Internet is not a bookshelf. Those who treat the Internet like a bookshelf miss the point about its impact, which is not just to replace existing objects and services.

In mid 2010 i searched Wikipedia for biographies of the main politicians of China and consistently found adulatory comments with virtually no reference to the role that those politicians (including Mao) played in blatant violations of human rights. Instead i found sentences such as "advocated China's peaceful development". In my research for my book on Silicon Valley i flipped through thousands of Wikipedia pages about companies and individuals: the vast majority were simply the equivalent of press releases worded according to the current business strategy of the company or according to the whims of the individual. In late 2010 the article on Feminism presented Mohammed (the founder of Islam) as the first major feminist in the history of the world. In february 2011 the article on detective fiction mentioned the medieval Arabian collection of stories "One Thousand and One Nights" as the first suspenseful book. Wikipedia pages on albums and films routinely describe them with a "Response from the critics was generally positive" comment, completely omitting the devastating reviews published by reliable critics. Obviously, these were all cases in which someone with a specific agenda was trying to influence the millions of people who rely on Wikipedia.

I started noticing a disturbing fact: the popularity of Wikipedia is de facto obliterating all the alternative sources that one could use to doublecheck Wikipedia articles. A search on any major topic routinely returns a Wikipedia page in the first two or three lines. The

other lines in the first page of results are almost inevitably commercial in nature. In order to find a scholarly page that can prove or disprove the Wikipedia page, one has to flip through several pages of results. Very few people make the effort. Therefore Wikipedia is rapidly becoming the only source of information about any major topic. Maybe this is acceptable for scientific topics (although i would still prefer that my Quantum Physics and Cellular Biology came from someone who has signed the article with her/his name and affiliation) but it is dangerous for topics that are "politicized" in nature. Then Wikipedia becomes the only source that millions of people access to find out what a politician, a government or a company has done. Worse: every topic can be "politicized" to some extent. I found references to the Bible and the Quran in articles about scientific topics. No traditional encyclopedia and no academic textbook in the free world would reference the Bible or the Quran to explain Quantum Mechanics or Cellular Biology. Precisely because it is edited by the general public, Wikipedia lends itself to a global politicization of every topic. It is an illusion that Wikipedians carry out "anonymous and collaborative editing": the very nature of Wikipedia encourages people to avoid collaboration and instead to leak ideological agendas into encyclopedia pages. The "collaboration" of which Wikipedia boasts is the fact that someone can retaliate to an opinionated or biased statement by removing or altering that statement and maybe inserting one that leans in the opposite direction; but a brawl is a very loose definition of "collaboration".

That danger is very visible in the rapid decline of quality. Like any corporation that has to hide its own shortcomings, Wikipedia boasts study after study that shows Wikipedia to be as accurate and more complete than the Encyclopedia Britannica. This is true only if one ignores semantics. In reality, there has never been and never will be a Britannica article that is simply the press release from a company or a doctored biography from a tyrannical government. If one considers the semantics, the gap between the accuracy of the

traditional encyclopedia and the inaccuracy of Wikipedia is rapidly increasing.

The evil is, obviously, not coming from the founder or the staff. It originates from the success itself of Wikipedia. According to a diagram from a 2011 presentation by Zack Exley that i attended, the number of senior (unpaid) Wikipedia editors rapidly reached 60,000 and has declined a bit during the Great Recession. That number, of course, does not tell the whole story. The meaningful number is the number of pages that an average unpaid editor has to maintain. In 2003 (just before the Wikipedia explosion) there were less than 200,000 articles and about 60,000 editors: on average three pages per senior editor. In 2010 the number of editors declined to 50,000 while the number of articles in English alone had increased to ten million (according to a diagram that is currently posted on the Wikipedia website (<http://en.wikipedia.org/wiki/File:EnwikipediaArt.PNG>): even assuming that all those 50,000 editors stick to Wikipedia's original philosophy (i'll say later why i don't believe it), that means 200 articles on average per editor.

Here is the bigger problem. When there were only a few thousand users, there was little interest from governments and corporations in what Wikipedia said. Now that there are millions of users and that the Wikipedia page is usually the first one presented by a search engine, the interest in determining what Wikipedia displays is enormous. There has been an undocumented explosion in the number of Wikipedia editors who are "paid" (either salaried or contracted) by governments, organizations, corporations and celebrities to twist the text of a Wikipedia article so that it represents the interest of that government, organization, corporation or celebrity.

When there were only a few thousand articles, it was relatively easy for the unpaid idealistic editors to control the content of Wikipedia. Now that there are millions of articles, it is simply impossible for those unpaid idealistic editors to control what the paid editors do. To make matters worse, Wikipedia covets the idea

that editors have to be anonymous: therefore there is no way for an unpaid idealistic editor to know if another editor is unpaid or paid. It's like those horror movies in which there is no way for a human to know whether she is surrounded by humans or zombies.

Like any corporation that has to hide its own shortcomings, Wikipedia boasts that "In the month of July 2006, Wikipedia grew by over 30,000,000 words". But that's precisely the problem. That's precisely what is frightening. Many of those 30 million words may be written by unprofessional, biased and sometimes paid "editors" whose interest in creating an encyclopedia is much lower than their interest in promoting a viewpoint or serving their employer. This leaves less than 50,000 unpaid idealistic Wikipedia editors to fight against an increasing number of editors paid by government agencies, ideological organizations, corporations and celebrities, not to mention the thousands of occasional uninformed amateurs who want to shout their opinion to the world.

Needless to say, a government agency, an ideological organization, a corporation or a celebrity has more resources at its disposal, and is much more determined, than a hapless unpaid Wikipedian. Therefore their version of the facts will eventually win. No wonder that an increasing number of pages simply displays what the subject of the page wants people to read. It is pointless for an idealistic editor to fight against it: the corporation or organization interested in that page has overwhelming resources to win the fight. There is no "brawl" over the content of those pages because it would be pointless. The net result is that Wikipedia is inevitably being hijacked by entities whose goal is not to spread knowledge but to spread propaganda.

Furthermore, several governments around the world block Wikipedia webpages. In the Middle East we were not able to access pages about Israel and Islam. In mainland China we could not access just about any page about history, including my own website www.scaruffi.com However, the free world can view the pages that have been doctored by the Chinese government and by Islamic religious groups. Therefore there is a one-way flow of mental

conditioning: "their" people cannot see our version of the facts, but we are increasingly exposed to "their" version of the facts as "they" become more and more active in editing Wikipedia pages. It is not difficult to predict who will win in the long run.

For government agencies, ideological organizations, corporations and celebrities Wikipedia has become a fantastic device to brainwash not only your own audience but all the people in the world.

Politically speaking, Wikipedia is de facto a force opposed to the change that social media foster. While Facebook and Twitter cannot be easily hijacked by authorities and corporations to brainwash people with distorted facts, Wikipedia can be and is being used precisely for that purpose by an increasing number of skilled and sinister "editors". Wikipedia can potentially become a force to stop change and promote repression, corruption, speculation and possibly genocide. Because they are so distributed and cannot be "edited", the voices expressed by Facebook and Twitter represent the voice of the people. The centralized Wikipedia, instead, is increasingly representing the voice of the oppressor; or, if you prefer, the oppressors are increasingly keen on appropriating Wikipedia.

In parallel, Wikipedia is having another detrimental effect on culture: it is sending out of business the only sources that we can use to verify Wikipedia's accuracy: the traditional encyclopedias. Compiling an encyclopedia is a colossal endeavor that requires the collective work of dozens of distinguished scholars. The cost for the publisher is enormous. In the age of Wikipedia no publisher is crazy enough to invest millions for an encyclopedia that will have to compete against the much bigger and absolutely free of charge Wikipedia. The age of encyclopedias that began in the Enlightenment is ending in the 21st century. In other words, the fact that Wikipedia is free has created a problem of historical proportions. Since no more encyclopedias will be produced, and any specialized website will be infinitely difficult to find using a search engine, society will have no way to determine if a Wikipedia

article is telling the truth or not. There will be no second source where one can doublecheck a statement, a date, a story, and let alone discuss the merits of who is represented on Wikipedia and who is not. Wikipedia is sending out of business the very sources that we use to determine Wikipedia's reliability and accuracy, the very sources that we used for centuries to determine the veracity of any statement. Wikipedia is not an encyclopedia, it is becoming a colossal accumulation of propaganda and gossip. The destruction of the traditional encyclopedia may send us back to the dark ages that followed the collapse of the Roman Empire.

P.S.

Things may be getting even more sinister as i write this book. Wikipedia's claim that anybody can edit an article is rapidly becoming an illusion: in reality, millions of IP addresses are banned from editing Wikipedia. A Stanford friend who added a link to a Wikipedia article (linking to this very article of mine) has never been able to edit articles again: Wikipedia displays an error message in which he is accused of "non constructive behavior". If this reminds you of totalitarian regimes, welcome to the world of Wikipedia. Wikipedia, by its own admission, keeps a detailed record of what every IP address in the world has written on which articles. And de facto Wikipedia bans from editing its pages those places (like libraries) that don't allow it to track down the identity of the person by the IP address. This is exactly what secret police like the KGB have always done in totalitarian regimes in which you are supposed to read (what they want you to read) but not to write (what you would like the world to read).

The usual objection to this comparison of mine is that Wikipedia editors are volunteers who do it just because they believe in the ideal. You'd be surprised how many members of the secret police in places like Nazi Germany, the Soviet Union and today's Iran were and are fanatical volunteers who believed in the ideal of their totalitarian state and were willing to work for free to fight the enemies of the state. The real enemy is often not the dictator in

charge but the fanatics who legitimize that dictator. Without those fanatical followers the totalitarian state would collapse.

Most users of Wikipedia have trouble accepting that Wikipedia is bad for humankind. They admit the limits and the potential harm, but would not want to erase it from the Web. My friend Federico Pistono, author of "Robots Will Steal Your Job, But That's OK" (2012), told me: "We just need to educate people on how to use it". My counter-suggestion: we should introduce more mistakes. It is important that the users of Wikipedia get "educated" to the idea that Wikipedia articles are typically biased articles written by whoever has more time and more money to continue editing them. In the interest of the truth, please change an article on the Nazi massacre of Jews in Poland so that "Warsaw" becomes "Acapulco" and "Hitler" becomes "Mickey Mouse". This way people will be aware that they cannot trust an anonymous Wikipedia article and they have to use other sources to doublecheck the content of Wikipedia articles. Sure: Wikipedia is useful to find out that Paris is the capital of France, and that the population of Nigeria is 173 million. It is very "useful" for many purposes. As long as we don't make excessive claims about its reliability: it is NOT an encyclopedia. At best, it is just a collection of advices given by amateurs to amateurs, just like reviews on Yelp and Amazon. Many television shows, documentaries and Internet videos have been useful to raise awareness about world events, but (hopefully) people know that those shows are run by comedians, entertainers and amateurs. Because Wikipedia articles are anonymous, people are routinely misled into thinking that they were written by top authorities more reliable than comedians and entertainers. In many cases that is not true. In fact, i don't know a single scholar who has contributed to a Wikipedia article.

How about a big banner on every Wikipedia article that warns "Disclaimer: None of the texts published here was provided or verified by a competent scholar"? Just like we warn people that cigarettes cause cancer?

Appendix: The Myth of Longevity

The new cult of digital immortality goes hand in hand with the widely publicized increases in life expectancy.

For centuries life expectancy (at older ages) rose very little and very slowly. What truly changed was infant mortality, that used to be very high. But towards the end of the 20th century life expectancy posted an impressive increase: according to the Human Mortality Database, in developed countries life expectancy at age 85 increased by only about one year between 1900 and 1960, but then increased by almost two years between 1960 and 1999. I call it the "100 curve": for citizens of developed countries the chance to live to 100 is now about 100 times higher than it was 100 years ago. In fact, if one projects the current trends according to the Human Mortality Database, most babies born since 2000 in developed countries will live to be 100.

James Vaupel, the founding director of the Max Planck Institute for Demographic Research, showed that the rate of increase in life expectancy is about 2.5 years per 10 years ("Demography", 2002). It means that every day our race's life expectancy increases by six hours. And Vaupel argues that life expectancy is likely to keep increasing.

These studies, however, often neglect facts of ordinary life. Since 1960 the conditions in which people live (especially urban people) have improved dramatically. For centuries people used to live with (and die because of) poor sanitation. The water supply of cities was chronically contaminated with sewage, garbage and carrion. Typhoid, dysentery and diarrhea were common. Outbreaks of smallpox, measles, polio, cholera, yellow fever, assorted plagues and even the flue killed millions before the invention of vaccines and the mandatory immunization programs of the last century. Before the 1960s polio would paralyze or kill over half a million people worldwide every year. Smallpox used to kill hundreds of thousands of Europeans annually (it was eradicated in 1979) and killed millions in the Americas after colonization. The World Health

Organization estimates that measles has killed about 200 million people worldwide over the last 150 years (but almost nobody in developed countries in recent decades). Cholera killed 200,000 people in the Philippines in 1902-04, 110,000 in the Ukraine in 1910 and millions in India in the century before World War I. The flu killed at least 25 million people worldwide in 1918, four million in 1957 and 750,000 in 1968. These causes of death virtually disappeared from the statistics of developed countries in the last half century. After 1960 diseases are generally replaced by earthquakes, floods and hurricanes (and man-made famines in communist countries) in the list of the mass killers. The big exceptions, namely tuberculosis (more than one million deaths a year), AIDS (almost two million deaths a year) and malaria (more than 700,000 deaths a year), are now mostly confined to the under-developed countries that are not included in the studies on life expectancy (the World Bank estimates that 99% of deaths due to these three diseases occur in underdeveloped countries).

Another major factor that contributed to extending life expectancy is affordable professional health care. Health care used to be the responsibility of the family before it shifted towards the state. The state can provide more scientific health care, but it is expensive. Professional health care became affordable after World War II thanks to universal health care programs: France (1945), Britain (1948), Sweden (1955), Japan (1961), Canada (1972), Australia (1974), Italy (1978), Spain (1986), South Korea (1989), etc. Among major developed countries Germany (1889) is the only one that offered universal health care before World War II (and the USA is the only one that still does not have one in place).

After improved health care and reduced infectious disease rates, the economist Dora Costa's "Causes of Improving Health and Longevity at Older Ages" (2005) lists "reduced occupational stress" and "improved nutritional intake" as the other major factors that determine longevity. However, work stress is increasing for women, as they ascend the corporate ladder, and data on diets (e.g., obesity) seem to point in the opposite direction: people quit

smoking, but now eat junk, and too much of it (and insert here your favorite rant against pesticides, hormone-raised meat and industrialized food in general).

Violent deaths have also decreased dramatically throughout the developed world: fewer and less bloody wars, and less violent crime. The rate of homicide deaths per 100,000 citizens is widely discussed in Steven Pinker's "The Better Angels of Our Nature" (2011). (Even in the USA where guns are widely available, and therefore violent crime kills exponentially more people than in Europe or Asia, the gun homicide rate decreased 49% between 1993 and 2013).

These factors certainly helped extend life expectancy in the developed world, but there is little improvement that they can still contribute going forward. In some cases one can even fear a regression. For example, no new classes of antibiotics have been introduced since 1987 whereas new pathogens are emerging every year, and existing bugs are developing resistance to current antibiotics. On the same day of March 2013 that a symposium in Australia predicted drugs to slow the ageing process within a decade so that people can live to 150 years the Chief Medical Officer for England, Dame Sally Davies, raised the alarm that antibiotics resistance may become a major killer in the near future. The Lancet, the British medical journal, estimated that in 2013 more than 58,000 babies died in India because they were born with bacterial infections that are resistant to known antibodies.

And we have to hope that the powers of the world will keep abstaining from using their nuclear weapons or we shall return to the statistics of the two world wars and possibly worse.

In 2013 the Centers for Disease Control and Prevention (CDC) calculated that every year in the USA at least two million people get ill because of antibiotic-resistant infections, and that at least 23,000 of them die. Drug-resistant tuberculosis killed an estimated 170,000 people in 2012.

The American Cancer Society calculated 1.6 million new cases of cancer and nearly 600,000 deaths in the USA in 2012, which

means that the number of cancer deaths in the USA has increased by 74% since 1970. The World Health Organization's "World Cancer Report 2014" estimated that cancer cases will increase by 70 percent over next 20 years.

The future promises more biomedical progress, and particularly therapies that may repair and reverse the causes of aging. This leads many to believe that human life can and will be extended dramatically, and maybe indefinitely.

However, health care has become too expensive for governments to continue footing the bill for the general population. Virtually every society in the developed world has been moving towards a larger base of elderly people and a smaller base of younger people who are supposed to pay for their health care. This equation is simply not sustainable. The professional health care that the average citizen receives may already have started to decline, and may continue to decline for a long time. It is just too expensive to keep the sick elderly alive forever for all the healthy youth who have to chip in. To compound the problem, statistics indicate that the number of people on disability programs is skyrocketing (14 million people in the USA in 2013, almost double the number of 15 years earlier). At the same time the tradition of domestic health care has largely been lost. You are on your own. This parallel development (unaffordable professional health care combined with the disappearance of domestic health care) is likely to reverse the longevity trend and lead to a worse (not better) chance of living a long life.

Furthermore, the rate of suicide has been increasing steadily in most developed societies, and, for whatever reason, it usually goes hand in hand with a decline in birth rate. Hence this might be an accelerating loop. The country with the oldest people is Japan. That is also one of the countries with the highest suicide rates of all, and most of the suicides are committed by elderly people. Getting very old does not make you very happy. In 2013 the Center for Disease Control (CDC) found that the suicide rate among middle-aged people in the USA had increased 28% in a decade (40% for white

people) and that since 2009 suicide had become the 10th leading cause of death in the country, overtaking car accidents.

As all countries reach the point of shrinking health care and accelerating suicide rates, life expectancy might actually start to decline for the first time in centuries.

Jeanne Louise Calment died at the age of 122 in 1997. Since then no person in the developed world (where you can verify the age) has died at an older age. Even if you believed the claims from various supercentennarians in developing countries (countries in which no document can prove the age of very old people), you could hardly credit their achievement on technological or medical progress since those supercentennarians lived all their lives with virtually no help from technology or medicine. In other words, the real numbers tell us that in almost 20 years nobody has reached the age that someone reached in 1997 with the possible exception of people who lived in underdeveloped countries. It takes a lot of imagination to infer from this fact that we are witnessing a trend towards longer lifespans.

There is also a shift in value perception at work. The idea that the only measure of a life is the number of years it lasted, that dying of old age is "better" than, say, dying in a car accident at a young age, is very much grounded in an old society driven by the survival instinct: survive at all costs for as long as possible. As the (unconscious) survival instinct is progressively replaced by (conscious) philosophical meditation in modern societies, more and more people will decide that dying at 86 is not necessarily better than dying at 85. In the near future people may care more about other factors than the sheer number of years they lived. The attachment to life and the desire to live as long as possible is largely instinctive and irrational. As generations become more and more rational about life, it may not sound so attractive to live long lives if one has to die anyway and be dead forever and be forgotten for the rest of eternity.

Then there are also many new habits that may contribute to creating a sicker species that will be more likely (not less likely) to die of diseases.

Most children in the developed world are now born to mothers (and fathers) aged 30 and older. As more parents delay childbearing, and the biological clock remains the same, the new generations are a veritable experiment (and sometimes literally a laboratory experiment). Fertility rates begin to decline gradually at age 30 and then decline exponentially, and to me that is nature's way of telling us when children should be born. In fact, babies born to older parents (as it is becoming common) have a higher risk of chromosome problems, as shown, for example, in a study led by Andrew Wyrobek and Brenda Eskenazi, "Advancing Age has Differential Effects on DNA Damage, Chromatin Integrity, Gene Mutations, and Chromosome Abnormalities in Sperm" (2006), and by a study led by Bronte Stone at the Reproductive Technology Laboratories, "Age thresholds for changes in semen parameters in men" (2013). Autism rates have risen 600 percent in the past two decades. While the age of the parents may not be the only cause, it is probably one significant cause. In Janie Shelton's and Irva Hertz-Picciotto's study "Independent and Dependent Contributions of Advanced Maternal and Paternal Ages to Autism Risk" (2010) the odds that a child would later be diagnosed with autism was 50% greater for a 40-year-old woman than for a woman between 25 and 29. To be fair, a study led by Mikko Myrskylä at the Max Planck Institute for Demographic Research in Germany, "Maternal Age and Offspring Adult Health" (2012), reassured many older mothers that education is the main factor determining the future health of their babies.

Nobody knows its causes and it is difficult to speculate on the effects, but several studies from European nations seem to show that the quality of sperm has deteriorated during the last half of a century. I doubt that this bodes well for the physical and mental health of our offspring. For example, a study led by the epidemiologist Joelle LeMoal ("Decline in Semen Concentration and

Morphology in a Sample of 26609 Men Close to General Population Between 1989 and 2005 in France", 2012) found that sperm concentration of young men decreased by nearly 30% in 17 years. If similar numbers showed up for environmental problems in a certain territory, we would immediately evacuate the place.

Last but not least, antibiotics, filtered water, cesarean-section childbirths and other environmental and behavioral aspects of modern life in developed countries greatly weaken the beneficial bacteria that constitute the physical majority of the cells of the human body. Vaccinations have been useful to prevent children from dying of horrible diseases, but now they are becoming mandatory for every possible disease (and for the mere possibility of a disease), thus creating weaker and weaker immune systems. Therefore health care itself (with its emphasis on vaccinations and antibiotics) may end up engineering weaker immune systems, which are much more likely to be defeated by unknown diseases than the unprotected immune system of our grandparents.

Personal mobility has greatly increased the chances that a deadly epidemic spreads worldwide killing millions of people.

Dana King's study "The Status of Baby Boomers' Health in the United States" (2013) seems to show that the "baby boomer" generation is less healthy than the previous generation. For example, twice as many baby boomers use a cane as did people of the same age in the previous generation. I personally have the feeling that the young people around me are less healthy than my generation was. Too many young people seem to have all sorts of physical problem and seem to get sick with every passing germ. They get food poisoning the moment they cross a border, and they start taking all sorts of pills in their 30s. I don't see how this bodes well for our race's longevity.

We now routinely eat genetically-manufactured food whose effects over the long time rest to be determined.

One fears that most of the gains in life expectancy may have already occurred, and now the challenge will be to preserve them. I

can't shake the feeling that we are building a weaker and weaker species while creating a more and more dangerous world.

Appendix: The Medium is the Brain

The World-wide Web has popularized the paradigm of navigating linked documents. This is an example of a process that introduces many distractions and inevitably reduces the depth of understanding (or, at least, increases the effort one has to make in order to stay focused). Generally speaking, the life of the individual who is permanently plugged into the network (hyperlink navigation, instant messages, live news) has a cost: the continuous shift of context and therefore of focus takes a cognitive toll on the brain.

Every time the brain has to reorient itself there is a cost in accessing long-time memory and organizing one's "thoughts". That brain gets trained for short attention spans. It is physically a different kind of brain from the brain that meditates and contemplates; from the brain that is capable of "deep understanding". The latter is trained by linear "texts" (or lectures or videos or whatever) that require the brain to remain on the same subject for a longer period of time and ransack long-term memory for all the appropriate resources to "understand" as much as possible. Brains that are trained to process linear texts comprehend more, remember more and, in my opinion, learn more, something already found in Erping Zhu's study "Hypermedia Interface Design" (1999). People who read linear text comprehend more, remember more, and learn more. Brains that are trained to switch focus all the time comprehend less, remember less and, possibly, learn less, as argued by Nicholas Carr in "The Shallows" (2010). This is due to the fact that it is "expensive" for the brain to transfer information from working memory to long-term memory (the "cognitive load"). Cognitive "overload" makes it difficult for the brain to decode and store information, and to create the appropriate links to pre-existing memories.

Guinevere Eden discussed how literacy reorganizes the brain at the physical level in "Development of Neural Mechanisms For Reading" (2003): reading and writing hijack a brain (so do other symbolic activities and art). Patricia Greenfield's study "Technology and Informal Education" (2009) shows that every medium develops some cognitive skills at the expense of others. Gary Small's "Patterns of Cerebral Activation during Internet Searching" (2009) proves how digital technology is rapidly and profoundly altering our brains. Betsy Sparrow's "Google Effects on Memory" (2011) shows how search engines change the way people use memory.

The medium that we use defines how the brain works. Ultimately, the medium physically changes our brain. The medium shapes the brain.

Every medium fosters some cognitive skills in the brain, but at the expense of others. There is a sort of zero sum of cognitive skills. A blind person improves smell and hearing. A videogame addict improves her visual-spatial skills but at the expense of other skills. The "focused" brain has skills that have been created by, for example, books, whereas the "switching" brain has skills that have been created by, for example, the Web.

The "switching" brain will lead to a more superficial society, in which brains are less and less capable of deep understanding. This is actually a process that has been going on for some centuries (if not millennia). At the time of Homer many people could memorize a lengthy poem. Before the invention of writing, brains had to memorize many more items than after the invention of writing. Before the invention of the specialist, people had to be experts in many fields of life, from carpentry to plumbing. After the invention of the society of specialists, we don't quite know how things work: we just know that by touching a switch or a lever something happens (a light comes on, a garage opens, a tv set turns on, water comes out of a faucet). The history of civilization is a history of reducing the amount of cognitive skills required to survive. Civilizations have constantly been refining the process of finding and using knowledge at the expense of the process of storing and understanding

knowledge. The Web-based society is simply a further step in this process, where navigating and multi-tasking prevail over deep understanding. We don't need to understand how things happen but just how to make things happen (e.g., if you want light, press a switch). Eventually, human brains may not be able to understand anything of the world that they "navigate" but will be able to do a lot more a lot faster in it.

This society of superficial brains will inevitably change the meaning of what is important. Science, literature and art were at the top of the hierarchy when deep understanding was important. Culture is not democratic at all. The academia decides what is more important and what is less important. In a society of superficial brains that don't need to understand much it is debatable whether a classic poem is still more important than a pulp novel. The elite-controlled hierarchy of knowledge becomes pointless in a world of superficial brains.

The switching brain works in fundamentally different ways and inevitably creates a fundamentally different society of brains. Literacy reorganizes the brain at the physical level: reading and writing hijack a brain; browsing and searching hijack a brain too. Here are some of the changes in the way the switching brain works.

The Web has so much information that one does not need intelligence anymore to solve a problem: most likely the solution can be found by navigating hyperlinked pages on the Web. The new way to solve a problem is not to concentrate on the nature of the problem, study the dynamics of the system and then logically infer what the solution could be. The new way is to search the Web for the solution posted by someone who knows it. At one point Artificial Intelligence was trying to build "expert systems" that would use knowledge and inference to find solutions. The Web makes the amount of knowledge virtually infinite and reduces the inference required by problem solving to just searching the knowledge for an adequate match. No mathematical logic needed. We are evolving towards a less and less intelligent way of solving problems, albeit

possibly a more and more effective way. The cognitive skill that we are losing is logical inference.

The combination of Web search and smartphones is also removing the need to think and argue about the truth of a statement: you can just "google" it and find the answer in a few seconds. There is no need to have a lengthy and emotional argument with a friend about who came first, the French or the USA revolution: just "google" it. Before the advent of the smartphone, one had to use all the brain's inferential skills and all the knowledge learned over a lifetime to guess the correct answer and to convince the audience. And the whole effort could easily lead to a wrong answer to be accepted by everybody. But that was a cognitive skill: rhetoric.

By the same token, there is no need to use a brain's orientation skills to find a place: just use the navigation system of the car or the smartphone. This removes the need to think and argue about whether to turn right or left. Before the advent of navigation systems, one had to use all the brain's inferential skills and all the knowledge learned over a lifetime to guess which way to go. And the whole effort could easily lead to picking a wrong direction. But that was a cognitive skill: orientation.

As our brain becomes more "superficial" it is likely that we also become more superficial in dealing with other individuals and with our world at large (family, friends, community, nation, our own life). One cognitive skill that may get lost in the age of "social networking" is precisely: socializing.

One skill that the switching brain is acquiring in place of the "focusing" skills is the skill of "posting" information. Before the Web only an elite was capable of producing content for the masses. The Web has created a large population of "prosumers", who are both passive consumers of content found on the Web and active producers of content for the Web (hats off to Alvin Toffler who coined the term in his 1980 book "The Third Wave", when the Internet was still an experiment). Social networking software, in particular, encourages people to post news about themselves, thus

creating a diary read by (potentially) millions of people. This is fostering a cognitive skill about "marketing" yourself to the world, about how to present your personality and your life to the others.

The simple act of browsing the Web constitutes a new cognitive skill. The browser is becoming de facto a new organ of the body, an organ used to explore the virtual universe of the Web, just like a hand or an eye is used to explore the physical universe. This organ is generating a new sense just like the hand created the sense of touch and the eye created the sense of sight. This new sense implies a new function in the brain just like any sense implies a corresponding function in the brain.

The switching brain must also be refining another skill that has been evolving over the last century: choice. Before the invention of cable television and the multiplication of channels the viewer had little choice on what to watch. For example, there were only a few evening news programs (in some countries only one on the national channel). The whole country was watching the same news at the same time. There was no need for searching and choosing the news. Cable television and now the Web have multiplied the possible sources of news and made them available around the clock. The "superficial" brain may not want to delve deeply into any particular event but probably needs to be much more skilled at searching and choosing the news. Choice is also involved in social networking systems to decide what is worth discussing, what is worth knowing and what is worth telling others about yourself.

On the other hand, it is not only that tools influence our being, but also that our being influences tools. The story is as much about how tools use our brains as about how our minds use tools. Often people end up using a tool in a way that is not the one it was designed for. This is particularly obvious in the case of software applications, but also in the case of many technologies that became runaway successes "despite" what the inventors originally intended for them. So much so that different people may use the same tool in different manners for different purposes (e.g., Facebook). We

express ourselves with the tools we have made as much as we see ourselves in the tools we have made.

The Web is the latest in a long series of new media that have shaped the human brain, starting with facial expression, language and writing. At each point some old skills were lost and some new skills were acquired. Your brain "is" the medium that shaped it. For better and for worse, you "are" the gadgets that you use.

Appendix: The Era of Objects

As we speculate on what will be the next stage of life's evolution, we underestimate the real innovation that has happened since the invention of intelligent life: objects. Life started building objects.

Life has not evolved much in the last ten thousand years, but objects have: there has been an explosive proliferation of objects.

We tend to focus on objects that mimic life (robots and the likes) as candidates for replacing life as we know it, and in the long term that might well be correct, but the objects that have truly multiplied and evolved at an astonishing rate are the ordinary static objects that populate our homes, our streets and our workplaces. There are objects virtually for anything.

When we look at life's evolution, we tend to look at how sophisticated the brain has become, but we tend to underestimate what that "sophisticated" brain is built to do: make more objects. The chimp's brain and the human brain are not all that different, and the behavior of the two species (eating, sleeping, sex, and perhaps even consciousness) are not all that different from the viewpoint of a (non-chimp and non-human) external observer, but the difference in terms of objects that they make is colossal. The real evolution of the brain is in terms of the objects it can build.

What the human race has truly accomplished is to turn a vast portion of the planet into objects: paved streets and sidewalks, buildings, cars, trains, appliances, clothes, furniture, kitchenware, etc.

Our lives revolve around objects. We work to buy a car or a home, and our work mostly consists in building or selling objects. We live to use them (usually in conjunction with other objects), to place them somewhere (usually inside other objects), to clean them (using other objects), etc.

The fundamental property of life is that it dies. Everything that was alive is now dead, except for those that are dying now. For us the planet is just a vast cemetery. For objects, instead, this planet is a vast factory of objects because, before dying, each of us builds or buys thousands of objects that will survive us and that will motivate future generations to build and buy more objects.

It is neither living beings nor genes nor memes (ideas) that evolve and drive evolution on this planet: it is objects. Objects have evolved far faster than life or ideas. The explosive proliferation of objects is the one thing that would be visible to anyone playing the last ten thousand years of history on Earth. Everything else (politics, economics, natural disasters, etc) pales in comparison to the evolution and proliferation of objects. The human body has not changed much in 200,000 years. Ideas have changed but slowly. Objects, instead, have completely changed and keep changing rapidly.

For example, what caused the collapse of the Soviet Union (and of communism in general) was neither the Pope nor Afghanistan: it was consumerism. Soviet citizens wanted goods in their stores, lots of goods. They actually liked many features of the communist society (they still do) but they wanted the proliferation of goods that democratic/capitalist societies offer. It was all about objects. Hence the Soviet Union collapsed because it dared challenge the domination of objects. For the same reason religions are declining everywhere: they are being replaced by philosophies of life that are more materialistic, i.e. that increase the evolution of objects.

Any system that challenges the absolute power of objects, or that doesn't contribute enough to the survival, proliferation and evolution of objects tends to lose. What benefits objects tends to succeed. Objects rule. Perhaps we are merely supposed to follow their

orders, and that's the only meaning of life. We get annihilated if we dare contradict objects.

You think that you are changing a car because you want a new car, but you can also see it the other way around: it is cars that want you to spend money that will go into making more and better cars.

In a sense, the consumer society is one stage in the evolution of objects, invented by objects in order to speed up their own evolution. Consumers are just the vehicle for objects to carry out their strategy of proliferation and domination.

Eventually objects will evolve into space stations and extraterrestrial colonies in order to expand outside this planet and begin the colonization of other parts of the universe in their quest to dominate all the matter that exists, until all matter in the universe will have been turned into objects by objects-creating beings like us (in clear defiance of the second law of Thermodynamics).

We are even turning our food into objects as we increasingly eat packaged food. The food system has changed more over the last 40 years than in the previous 40,000 years.

Shoes, refrigerators, watches and underwear are the real protagonists of history. Everything else is just a footnote to their odyssey.

Appendix: I Ran Out of Space

You can read more essays that didn't fit in this book at www.scaruffi.com/essays.html

A Timeline of Neuroscience

1590: Rudolph Goeckel's "Psychologia" introduces the word "psychology" for the discipline that studies the soul

1649: Pierre Gassendi's "Syntagma philosophiae Epicuri" argues that beasts have a cognitive life of their own, just inferior to humans

1664: Rene Descartes' "Treatise of Man" argues that the pineal gland is the main seat of consciousness (Great Minds Series):

1664: Thomas Willis' "Cerebral Anatomy" (1664) describes the different structures in the brain and coins the word "neurology"

1741: Emanuel Swedenborg's "The Economy of the Animate Kingdom" discusses cortical localization in the brain

1771: Luigi Galvani discovers that nerve cells are conductors of electricity

1796: Franz-Joseph Gall begins lecturing on phrenology, holding that mental faculties are localized in specific brain regions (of which 19 are shared with animals and 8 are exclusive to humans)

1824: Pierre Flourens' "Phrenology Examined" discredits Gall

1825: Jean-Baptiste Bouillaud's "Clinical and Physiological Treatise upon Encephalitis" describes patients who suffered brain lesions and lost their speech ability

1836: Marc Dax's "Lesions of the Left Half of the Brain Coincident With the Forgetting of the Signs of Thought" notes that aphasic patients (incapable of speaking) have sustained damage to the left side of the brain

1861: Paul Broca's "Loss of Speech, Chronic Softening and Partial Destruction of the Anterior Left Lobe of the Brain" single-handedly resurrects the theory of cortical localization of function

1865: Paul Broca's "Localization of Speech in the Third Left Frontal Convolution" suggests that the location of speech must be in the left hemisphere

1868: John Hughlings Jackson's "Notes on the Physiology and Pathology of the Nervous System" reports how damage to the right hemisphere impairs spatial abilities

1870: Eduard Hitzig and Gustav Fritsch discover the location of the motor functions in the brain

1873: Jean-Martin Charcot's "Lectures on the Diseases of the Nervous System" describes the neural origins of multiple sclerosis

1873: Camillo Golgi's "On the Structure of the Brain Grey Matter" describes the body of the nerve cell with a single axon and several dendrites

1874: Karl Wernicke determines that sensory aphasia (a loss of linguistic skills) is related to damage to the left temporal lobe

1874: Charles-Edouard Brown-Sequard's "Dual Character of the Brain" argues that education does not adequately target the right hemisphere

1876: John Hughlings Jackson discovers that loss of spatial skills is related to damage to the right hemisphere

1876: David Ferrier's "The Functions of the Brain" provides a map of brain regions specialized in motor, sensory and association functions

1890: Wilhelm His coins the word "dendrite"

1891: Santiago Ramon y Cajal proves that the nerve cell (the neuron) is the elementary unit of processing in the brain, receiving inputs from other neurons via the dendrites and sending its output to other neurons via the axon

1891: Wilhelm von Waldeyer coins the term "neuron" while discussing Santiago Ramon y Cajal's theory

1896: Albrecht von Kolliker coins the word "axon"

1897: Charles Sherrington coins the word "synapse"

1901: Charles Sherrington maps the motor cortex of apes

1903: Alfred Binet's "intelligent quotient" (IQ) test

1905: Keith Lucas demonstrates that below a certain threshold of stimulation a nerve does not respond to a stimulus and, once the threshold is reached, the nerve continues to respond by the same fixed amount no matter how strong the stimulus is

1906: Charles Sherrington's "The Integrative Action of the Nervous System" argues that the cerebral cortex is the center of integration for cognitive life

1911: Edward Thorndike's connectionism (the mind is a network of connections and learning occurs when elements are connected)

1921: Otto Loewi demonstrated chemical transmission of nerve impulses, proving that nerves can excite muscles via chemical reactions (notably acetylcholine) and not just electricity

1924: Hans Berger records electrical waves from the human brain, the first electroencephalograms

1924: Konstantin Bykov, performing split-brain experiments on dogs, discovers that severing the corpus callosum disables communications between the two brain hemispheres

1924: Hans Berger records electrical waves from the human brain, the first electroencephalograms

1925: Edgar Adrian shows that the message from one neuron to another neuron is conveyed by changes in the frequency of the discharge, the first clue on how sensory information might be coded in the neural system

1928: Otfried Foerster stimulates the brain of patients during surgery with electric probes

1933: Henry Dale coins the terms "adrenergic" and "cholinergic" to describe the nerves releasing the two fundamental classes of neurotransmitters, the adrenaline-like one and acetylcholine

1935: Wilder Penfield explains how to stimulate the brain of epileptic patients with electrical probes ("Epilepsy and Surgical Therapy")

1936: Jean Piaget's "The Origins of Intelligence in Children"

1940: Willian Van Wagenen performs "split brain" surgery to control epileptic seizures

1949: Donald Hebb's cell assemblies (selective strengthening or inhibition of synapses causes the brain to organize itself into regions of self-reinforcing neurons - the strength of a connection depends on how often it is used)

1951: Roger Sperry's "chemoaffinity theory" of synapse formation explains how the nervous system organizes itself during embryonic development via a genetically-determined chemical matching program

1952: Paul Maclean discovers the "limbic system"

1953: John Eccles' "The Neurophysiological Basis of Mind" describes excitatory and inhibitory potentials, the two fundamental changes that occur in neurons

1953: Roger Sperry and Ronald Meyers study the "split brain" and discover that the two hemispheres are specialized in different tasks

1953: Eugene Aserinsky discovers "rapid eye movement" (REM) sleep that corresponds with periods of dreaming

1954: Rita Levi-Montalcini discover nerve-growth factors that help to develop the nervous system, thus proving Sperry's chemoaffinity theory

1957: Vernon Mountcastle discovers the modular organization of the brain (vertical columns)

1959: Michel Jouvet discovers that REM sleep originates in the pons

1962: David Kuhl invents SPECT (single photon emission computer tomography)

1964: John Young proposes a "selectionist" theory of the brain (learning is the result of the elimination of neural connections)

1964: Paul Maclean's triune brain: three layers, each layer corresponding to a different stage of evolution

1964: Lueder Deecke and Hans-Helmut Kornhuber discover an unconscious electrical phenomenon in the brain, the Bereitschaftspotential (readiness potential)

1964: Benjamin Libet discovers that the readiness potential precedes conscious awareness by about half a second

1968: Niels Jerne's selectionist model of the brain (mental life a continuous process of environmental selection of concepts in our brain - the environment selects our thoughts)

1972: Raymond Damadian builds the world's first Magnetic Resonance Imaging (MRI) machine

1973: Edward Hoffman and Michael Phelps create the first PET (positron emission tomography) scans that allow scientists to map brain function

1972: Jonathan Winson discovers a correlation between the theta rhythm of dreaming and long-term memory

1972: Godfrey Hounsfield and Allan Cormack invent computed tomography scanning or CAT-scanning

1977: Allan Hobson's theory of dreaming

1978: Gerald Edelman's theory of neuronal group selection or "Neural Darwinism"

1985: Michael Gazzaniga's "interpreter" (a module located in the left brain interprets the actions of the other modules and provides explanations for our behavior)

1989: Wolf Singer and Christof Koch discover that at, any given moment, very large number of neurons oscillate in synchrony and one pattern is amplified into a dominant 40 Hz oscillation (gamma synchronization)

1990: Seiji Ogawa's "functional MRI" measures brain activity based on blood flow

1994: Vilayanur Ramachandran proves the plasticity of the adult human brain

1996: Giacomo Rizzolatti discovers that the brain uses "mirror" neurons to represent what others are doing

1996: Rodolfo Llinas: Neurons are always active endlessly producing a repertory of possible actions, and the circumstances "select" which specific action is enacted

1997: Japan opens the Brain Science Institute near Tokyo

2009: The USA launches the Human Connectome Project to map the human brain

2012: Mark Mayford stores a mouse's memory of a familiar place on a microchip

2013: The European Union launches the Human Brain Project to computer-simulate the human brain

A Timeline of Artificial Intelligence

- 1935: Alonzo Church proves the undecidability of first order logic
- 1936: Alan Turing's Universal Machine ("On computable numbers, with an application to the Entscheidungsproblem")
- 1936: Alonzo Church's Lambda calculus
- 1941: Konrad Zuse's programmable electronic computer
- 1943: "Behavior, Purpose and Teleology" co-written by mathematician Norbert Wiener, physiologist Arturo Rosenblueth and engineer Julian Bigelow
- 1943: Kenneth Craik's "The Nature of Explanation"
- 1943: Warren McCulloch's and Walter Pitts' binary neuron ("A Logical Calculus of the Ideas Immanent in Nervous Activity")
- 1945: John Von Neumann designs a computer that holds its own instructions, the "stored-program architecture"
- 1946: The ENIAC, the first Turing-complete computer
- 1946: The first Macy Conference on Cybernetics
- 1947: John Von Neumann's self-reproducing automata
- 1948: Alan Turing's "Intelligent Machinery"
- 1948: Norbert Wiener's "Cybernetics"
- 1949: Leon Dostert founds Georgetown University's Institute of Languages and Linguistics
- 1949: William Grey-Walter's Elmer and Elsie robots
- 1950: Alan Turing's "Computing Machinery and Intelligence" (the "Turing Test")
- 1951: Claude Shannon's maze-solving robots ("electronic rats")
- 1951: Karl Lashley's "The problem of serial order in behavior"
- 1952: First International Conference on Machine Translation organized by Yehoshua Bar-Hillel
- 1952: Ross Ashby's "Design for a Brain"
- 1954: Demonstration of a machine-translation system by Leon Dostert's team at Georgetown University and Cuthbert Hurd's team at IBM, possibly the first non-numerical application of a digital computer
- 1956: Allen Newell and Herbert Simon demonstrate the "Logic Theorist"
- 1956: Dartmouth conference on Artificial Intelligence
- 1957: Frank Rosenblatt's Perceptron
- 1957: Newell & Simon's "General Problem Solver"
- 1957: Noam Chomsky's "Syntactic Structures" (transformational grammar)
- 1958: John McCarthy's LISP programming language

- 1958: Oliver Selfridge's Pandemonium
- 1959: Arthur Samuel's Checkers, the world's first self-learning program
- 1959: John McCarthy and Marvin Minsky found the Artificial Intelligence Lab at the MIT
- 1959: John McCarthy's "Programs with Common Sense" (1949) focuses on knowledge representation
- 1959: Noam Chomsky's review of a book by Skinner ends the domination of behaviorism and resurrects cognitivism
- 1959: Yehoshua Bar-Hillel's "proof" that machine translation is impossible
- 1960: Bernard Widrow's and Ted Hoff's Adaline ((Adaptive Linear Neuron or later Adaptive Linear Element) that uses the Delta Rule for neural networks
- 1960: Hilary Putnam's Computational Functionalism
- 1962: Joseph Engelberger deploys the industrial robot Unimate at General Motors
- 1963 Irving John Good (Isidore Jacob Gudak) speculates about "ultraintelligent machines" (the "singularity")
- 1963 John McCarthy moves to Stanford and founds the Stanford Artificial Intelligence Laboratory (SAIL)
- 1964: IBM's "Shoebbox" for speech recognition
- 1965: Ed Feigenbaum's Dendral expert system
- 1965: Lotfi Zadeh's Fuzzy Logic
- 1966: Joe Weizenbaum's Eliza
- 1966: Ross Quillian's semantic networks
- 1967: Barbara Hayes-Roth's Hearsay speech recognition system
- 1967: Charles Fillmore's Case Frame Grammar
- 1967: Leonard Baum's team develops the Hidden Markov Model
- 1968: Glenn Shafer's and Stuart Dempster's "Theory of Evidence"
- 1968: Peter Toma founds Systran to commercialize machine-translation systems
- 1969: First International Joint Conference on Artificial Intelligence (IJCAI) at Stanford
- 1969: Marvin Minsky & Samuel Papert's "Perceptrons" kill neural networks
- 1969: Roger Schank's Conceptual Dependency Theory for natural language processing
- 1969: Stanford Research Institute's Shakey the Robot
- 1970: Albert Uttley's Informon for adaptive pattern recognition
- 1970: William Woods' Augmented Transition Network (ATN) for natural language processing
- 1971: Ingo Rechenberg publishes his thesis "Evolution Strategies", a set of optimization methods for evolutionary computation

1971: Richard Fikes' and Nils Nilsson's STRIPS
 1972: Alain Colmerauer's PROLOG programming language
 1972: Bruce Buchanan's MYCIN
 1972: Hubert Dreyfus's "What Computers Can't Do"
 1972: Terry Winograd's Shrdlu
 1973: "Artificial Intelligence: A General Survey" by James Lighthill criticizes Artificial Intelligence for over-promising
 1974: Marvin Minsky's Frame
 1974: Paul Werbos' Backpropagation algorithm for neural networks
 1975: John Holland's genetic algorithms
 1975: Roger Schank's Script
 1976: Doug Lenat's AM
 1976: Richard Laing's paradigm of self-replication by self-inspection
 1979: Cordell Green's system for automatic programming
 1979: David Marr's theory of vision
 1979: Drew McDermott's non-monotonic logic
 1979: William Clancey's Guidon
 1980: Intellicorp, the first major start-up for Artificial Intelligence
 1980: John McDermott's Xcon
 1980: John Searle publishes the article "'Minds, Brains, and Programs'" on the "Chinese Room" that attacks Artificial Intelligence
 1980: Kunihiko Fukushima's Convolutional Neural Networks
 1980: McCarthy's Circumscription
 1981: Danny Hillis' Connection Machine
 1982: Japan's Fifth Generation Computer Systems project
 1982: John Hopfield describes a new generation of neural networks, based on a simulation of annealing
 1982: Judea Pearl's "Bayesian networks"
 1982: Teuvo Kohonen's Self-Organized Maps (SOM) for unsupervised learning
 1982: The Canadian Institute for Advanced Research (CIFAR) establishes Artificial Intelligence and Robotics as its very first program
 1983: Geoffrey Hinton's and Terry Sejnowski's Boltzmann machine for unsupervised learning
 1983: John Laird and Paul Rosenbloom's SOAR
 1984: Valentino Breitenberg's "Vehicles"
 1986: David Rumelhart's "Parallel Distributed Processing" rediscovers Werbos' backpropagation algorithm
 1986: Paul Smolensky's Restricted Boltzmann machine
 1987: Chris Langton coins the term "Artificial Life"
 1987: Hinton moves to the Canadian Institute for Advanced Research (CIFAR)

- 1987: Marvin Minsky's "Society of Mind"
- 1987: Rodney Brooks' robots
- 1987: Stephen Grossberg's Adaptive Resonance Theory (ART) for unsupervised learning
- 1988: Fred Jelinek's team at IBM publishes "A statistical approach to language translation"
- 1988: Hilary Putnam: "Has artificial intelligence taught us anything of importance about the mind?"
- 1988: Philip Agre builds the first "Heideggerian AI", Pengi, a system that plays the arcade videogame Pengo
- 1990: Carver Mead describes a neuromorphic processor
- 1992: Thomas Ray develops "Tierra", a virtual world
- 1994: The first "Toward a Science of Consciousness" conference in Tucson, Arizona
- 1995: Geoffrey Hinton's Helmholtz machine
- 1996: David Field & Bruno Olshausen's sparse coding
- 1997: IBM's "Deep Blue" chess machine beats the world's chess champion, Garry Kasparov
- 1998: Two Stanford students, Larry Page and Russian-born Sergey Brin, launch the search engine Google
- 1998: Yann LeCun's second generation Convolutional Neural Networks
- 2000: Cynthia Breazeal's emotional robot, "Kismet"
- 2000: Seth Lloyd's "Ultimate physical limits to computation"
- 2001: Juyang Weng's "Autonomous mental development by robots and animals"
- 2001: Nikolaus Hansen introduces the evolution strategy called "Covariance Matrix Adaptation" (CMA) for numerical optimization of non-linear problems
- 2003: Hiroshi Ishiguro's Actroid, a robot that looks like a young woman
- 2003: Jackrit Suthakorn and Gregory Chirikjian at Johns Hopkins University build an autonomous self-replicating robot
- 2003: Tai-Sing Lee's "Hierarchical Bayesian inference in the visual cortex"
- 2004: Ipke Wachsmuth's conversational agent "Max"
- 2004: Mark Tilden's biomorphic robot Robosapien
- 2005: Andrew Ng at Stanford launches the STAIR project (Stanford Artificial Intelligence Robot)
- 2005: Boston Dynamics' quadruped robot "BigDog"
- 2005: Hod Lipson's "self-assembling machine" at Cornell University
- 2005: Honda's humanoid robot "Asimo"
- 2006: Geoffrey Hinton's Deep Belief Networks (a fast learning algorithm for restricted Boltzmann machines)

- 2006: Osamu Hasegawa's Self-Organising Incremental Neural Network (SOINN), a self-replicating neural network for unsupervised learning
- 2007: Yeshua Bengio's Stacked Auto-Encoders
- 2008: A 3D Printer builds a copy of itself at the University of Bath
- 2008: Adrian Bowyer's 3D Printer builds a copy of itself at the University of Bath
- 2008: Cynthia Breazeal's team at the MIT's Media Lab unveils Nexi, the first mobile-dexterous-social (MDS) robot
- 2008: Dharmendra Modha at IBM launches a project to build a neuromorphic processor
- 2010: Lola Canamero's Nao, a robot that can show its emotions
- 2010: Quoc Le's "Tiled Convolutional Networks"
- 2010: The New York stock market is shut down after algorithmic trading has wiped out a trillion dollars within a few seconds.
- 2011: IBM's Watson debuts on a tv show
- 2011: Nick D'Aloisio releases the summarizing tool Trimit (later Summly) for smartphones
- 2011: Osamu Hasegawa's SOINN-based robot that learns functions it was not programmed to do
- 2012: Rodney Brooks' hand programmable robot "Baxter"
- 2013: John Romanishin, Kyle Gilpin and Daniela Rus' "M-blocks" at MIT
- 2014: Alex Graves, Greg Wayne and Ivo Danihelka publish a paper on "Neural Turing Machines"
- 2014: Jason Weston, Sumit Chopra and Antoine Bordes publish a paper on "Memory Networks"
- 2014: Li Fei-Fei's computer vision algorithm that can describe photos ("Deep Visual-Semantic Alignments for Generating Image Descriptions", 2014)
- 2014: Vladimir Veselov's and Eugene Demchenko's program Eugene Goostman, which simulates a 13-year-old Ukrainian boy, passes the Turing test at the Royal Society in London

Readings on the Singularity

- Brynjolfsson, Erik & McAfee, Andrew: "Race Against the Machine" (2012)
- Bostrom, Nick: "Are we living in a simulation?" (2003)
- Bostrom, Nick: "Superintelligence" (2014)
- Chalmers, David: "The Singularity - A Philosophical Analysis" (2010)
- Eden, Amnon & others: "The Singularity Hypotheses" (2013)
- Good, Irving John: "Speculations Concerning the First Ultraintelligent Machine" (1964)
- Kurzweil, Ray: "The Age of Intelligent Machines" (1990)
- Lanier, Jaron: "The Singularity Is Just a Religion for Digital Geeks" (2010)
- Moravec, Hans: "Mind Children: The Future of Robot and Human Intelligence" (1988)
- Muehlhauser, Luke: "When Will AI Be Created?" (2013)
- Scaruffi, Piero: "The Nature of Consciousness" (2006)
- Vinge, Vernon: "The Singularity" (1993)
- IEEE Spectrum's "Special Report: The Singularity" (2008)
- The Journal of Consciousness Studies's Volume 19/7-8 (2012):
- Corabi & Schneider: "The metaphysics of uploading"
 - Dainton, Barry: "On singularities and simulations"
 - Hutter, Marcus: "Can intelligence explode?"
 - Steinhart, Eric: "The singularity: Beyond philosophy of mind"

Alphabetical Index

- Anderson, William : 59
 Arakawa, Akio : 40
 Armstrong, Neil : 53
 Asimov, Isaac : 61
 Avery, Oswald : 59
 Bayes, Thomas : 21, 23, 140
 Bengio, Yeshua : 23
 Bettencourt, Luis : 58
 Bigelow, Julian : 10
 Binnig, Gerd : 145
 Bostrom, Nick : 135, 138
 Bratton, Benjamin : 61
 Breazeal, Cynthia : 23
 Breitenberg, Valentino : 20
 Brenner, Sydney : 58, 113
 Brock, David : 143
 Brooks, Rodney : 20, 23
 Brynjolfsson, Erik : 86
 Buchanan, Bruce : 12
 Calment, JeanneLouise : 179
 Camping, Harold : 9
 Carr, Nicholas : 182
 Carter, Brandon : 142
 Chamberlain, Fred : 114
 Chapin, John : 114
 Charney, Jule : 40
 Chirikjian, Gregory : 13
 Chomsky, Noam : 11
 Cooper, Evan : 114
 Costa, Dora : 176
 Cowen, Tyler : 154
 Cox, Earl : 112
 Crick, Francis : 59
 Dancer, John : 120
 Davies, Sally : 177
 Deamer, David : 15
 Debussy, Claude : 51
 Delgado, Jose : 114
 Deutsch, David : 67
 Dickmanns, Ernst : 25
 Diffenbaugh, Noah : 156
 Doyle, Richard : 112
 Drake, Edwin : 136
 Dubinsky, Donna : 22
 Eden, Guinevere : 183
 Ehrlich, Paul : 155
 Eigler, Don : 145
 Einstein, Albert : 47
 Ellison, Larry : 96
 Engelbart, Doug : 99
 Esfandiary, Fereidoun : 112
 Eskenazi, Brenda : 180
 Ettinger, Robert : 114
 Exley, Zack : 170
 Feigenbaum, Ed : 11
 Field, Chris : 155, 159
 Fukushima, Kunihiko : 22
 Gagarin, Yuri : 53
 George, Dileep : 22
 Gibbs, JosiahWillard : 21
 Gibson, William : 112
 Gilpin, Kyle : 35
 Giorgi, Filippo : 156
 Goedel, Kurt : 14, 68
 Good, Jack : 85
 Gopnik, Alison : 18, 82
 Grantham, Jeremy : 60, 158
 Green, Cordell : 14
 Grenander, Ulf : 21
 GreyWalter, William : 15
 Hammerbacher, Jeff : 97, 154
 Hansen, Nikolaus : 24
 Hanson, Robin : 160
 Hasegawa, Osamu : 23
 Hawkins, Jeff : 22
 Hayes, Barbara -Roth: 13

- Hegel, Georg : 69
 Heine, Heinrich : 118
 Heinrich, Andreas : 145
 Hertz, Irva -Picciotto: 180
 Hilbert, David : 7, 14
 Hinton, Geoffrey : 21, 22
 Holland, John : 13
 Hopfield, John : 20, 140
 Indiveri, Giacomo : 128
 Itskov, Dmitry : 113
 Jelinek, Fred : 29
 Kasparov, Garry : 35
 Kasthuri, Narayanan : 114
 Kaufman, Steve : 106
 King, Dana : 181
 Knuth, Donald : 140
 Koyama, Ken : 161
 Kramnik, Vladimir : 35
 Kurzweil, Ray : 6, 7, 24, 113
 Laclanche, Georges : 136
 Laing, Richard : 13
 Lancaster, Madeline : 137
 LeCun, Yann : 23
 Lee, Tai-Sing: 22
 LeMoal, Joelle : 180
 Leslie, John : 142
 Lichtman, Jeff : 114
 Lorenz, Edward : 40
 Marmer, Max : 154
 Martin, George : 112
 McAfee, Andrew : 86
 McCarthy, John : 10, 99, 130
 McCulloch, Warren : 126
 McDermott, John : 12
 McGinn, Colin : 66, 67
 McKay, Chris : 25
 Mead, Carver : 14, 128
 Meier, Karlheinz : 128
 Minsky, Marvin : 12
 Mitchell, George : 161
 Modha, Dharmendra : 128
 Moravec, Hans : 86
 More, Max : 112
 Morrow, Tom : 112
 Mullis, Kary : 59
 Mumford, David : 21
 Musk, Elon : 74
 Myrskylä, Mikko : 180
 Nash, Graham : 29
 Newell, Allen : 10
 Ng, Andrew : 26
 Nicoletis, Miguel : 114
 Nietzsche, Friedrich : 7
 Nunez, Paul : 37, 58
 Ohno, Taiichi : 88
 Packard, Vance : 96
 Paller, Alan : 117
 Papert, Samuel : 12
 Pearl, Judea : 21
 Picasso, Pablo : 146
 Pinker, Steven : 177
 Pistono, Federico : 174
 Pitts, Walter : 126
 Plante, Gaston : 136
 Pontin, Jason : 155
 Quillian, Ross : 13
 Reagan, Ronald : 30
 Rechenberg, Ingo : 13
 Rice, Henry : 131
 Richardson, Lewis : 39
 Rohrer, Heinrich : 145
 Romanishin, John : 35
 Rose, Charlie : 158
 Rosenblatt, Frank : 12
 Rosenblueth, Arturo : 10
 Rothblatt, Martine : 112
 Rus, Daniela : 35
 Russell, Bertrand : 7, 31, 73
 Ryle, Gilbert : 39
 Salisbury, Kenneth : 129
 Sandberg, Anders : 135
 Sanger, Frederick : 59
 Sarma, Sanjay : 143
 Schank, Roger : 13

- Schoenberg, Arnold : 51
Schopenhauer, Arthur : 7
Searle, Rick : 155
Sejnowski, Terry : 21
Selfridge, Oliver : 12
Shannon, Claude : 131
Shelton, Janie : 180
Simon, Herbert : 10, 33
Small, Gary : 183
Smith, Hamilton : 59
Smith, Hedrick : 88
Smolensky, Paul : 21
Snowden, Edward : 97
Socrates: 118
Sparrow, Betsy : 183
Sporns, Olaf : 113
Srinivasa, Narayan : 128
Stephenson, Neal : 154
Stewart, Robin : 40
Stravinsky, Igor : 51
Suthakorn, Jackrit : 13
Tegmark, Max : 31
Thiel, Peter : 154
Toffler, Alvin : 185
Toma, Peter : 13
Turing, Alan : 7, 9, 10, 14, 56, 68, 131
Varese, Edgar : 51
Vaupel, James : 175
Venter, Craig : 59
Von, John Neumann: 39
VonFoerster, Heinz : 59
VonNeumann, John : 128, 131, 144
Wachsmuth, Ipke : 23
Watson, James : 59
Watt, James : 49, 136
Weiser, Mark : 143
Weizenbaum, Joe : 13
West, Geoffrey : 58
White, John : 113
Wiener, Norbert : 10, 131
Wilkinson, John : 49
Wilmuth, Ian : 59
Winograd, Terry : 13, 27
Wittgenstein, Ludwig : 73
Wyrobek, Andrew : 180
Zhu, Erping : 182
Zuse, Konrad : 48