

*This informational blog was written by one of my best friends. He's a mathematician, scientist, and involved in building out the global infrastructure for the internet we all use today. In my high school circle of friends he was considered our "Spock". He's responding to an interesting blog I sent his way on [Artificial Intelligence](#) (AI) that was stirred up from my post CES writing on IBM's initiative on AI. It's a bit lengthy so you may want to download the pdf format for reading while on a flight or other casual moments in your life. I think you'll find it quite stimulating.*

## **Artificial Intelligence and Grandkids**

Far be it from me to downplay the risks of linear thinking in an exponential world; however, in my opinion, the arguments presented for Artificial General Intelligence (AGI) and Artificial Super Intelligence (ASI) are overly aggressive. I suspect it is much more likely that something spectacular (not predicting whether it would be a positive or a negative) will occur on the biological front, versus solely on the technological one; perhaps a blend of both.

As a reference, here are the definitions from the article:

**AI Caliber 1) Artificial Narrow Intelligence (ANI):** Sometimes referred to as *Weak AI*, Artificial Narrow Intelligence is AI that specializes in *one* area. There's AI that can beat the world chess champion in chess, but that's the only thing it does. Ask it to figure out a better way to store data on a hard drive, and it'll look at you blankly.

**AI Caliber 2) Artificial General Intelligence (AGI):** Sometimes referred to as *Strong AI*, or *Human-Level AI*, Artificial General Intelligence refers to a computer that is as smart as a human *across the board*—a machine that can perform any intellectual task that a human being can. Creating AGI is a *much* harder task than creating ANI, and we're yet to do it. Professor Linda Gottfredson describes intelligence as "a very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly, and learn from experience." AGI would be able to do all of those things as easily as you can.

**AI Caliber 3) Artificial Superintelligence (ASI):** Oxford philosopher and leading AI thinker Nick Bostrom [defines](#) superintelligence as "an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills." Artificial Superintelligence ranges from a computer that's just a little smarter than a human to one that's trillions of times smarter—across the board. ASI is the reason the topic of AI is such a spicy meatball and why the words "immortality" and "extinction" will both appear in these posts multiple times.

The very steep "S" curves up-and-to-the-right are, indeed, representative of what has occurred within relatively narrow technology areas; but, these can't simply be extrapolated onto more general topics that happen to make use of one or more of these technologies.

Consider the timeline proposed in the article for AGI in 25 years (roughly 2040) and ASI in 45 years (2060). Now, imagine us (you & me) spanning that AGI timeframe backwards and forwards (i.e. 25 years ago and 25 years ahead) which makes us 35-85 years of age, roughly; and it's not unreasonable to think that we might well see all of that. That's a 50 year window. On the ASI timeline, plus or minus 45 years, that 90 year window extends our age from 15 to 105. We probably won't see all of that, but we will see most of it. That takes us from 1970 thru to 2060. So, what's my point? For now, just remember AGI 2040 & ASI 2060.

First, as an example, let's look at commercial aviation; a fairly technological field, that leverages computer technologies, material sciences, meteorological advances, robotics, fiber optics, satellites, aerospace developments, etc, etc.

The Wright Brothers did their thing just after the turn of the century, the first commercial flight was in 1920, jet engines were introduced in 1952, the SSTs (aka Concorde) came online in 1976 (and were retired in 2003). Incremental improvements in range, capacity, conveniences (e.g. phones onboard, and now WiFi), fuel economies, lighter/stronger materials, fly-by-wire versus hydraulics, navigation, fewer cockpit crew to name a few have certainly been made in the 95 years since commercial flight began, with most of them occurring in the last 25 years or so. But, it has taken 95 years to get here, and there really isn't anything significant on the horizon to warrant a radical change in expectations. There are lots of reasons for this, including safety concerns for human life; but, commercial viability is just as critical a factor, if not the ultimate controlling variable. The SSTs failed because they couldn't turn a profit, so the return on investment was perpetually negative. It takes billions of R&D dollars and several years to develop a new jet engine that is only marginally more effective than its predecessor; there are no orders of magnitude improvements, which is the definition of exponential. Therefore, the payback period for this investment is measured in decades (like 30 years), so the overall project" life, from inception through retirement, is 50+ years. With that reality as a backdrop, there is no incentive to pursue anything more dramatic; nothing that would yield in 25 years the type of quantum leap required to achieve AGI, let alone ASI. In fact, there is an incredible disincentive at work to inhibit anything that might disrupt the commercial viability of the "incremental, proven, technology" path. (If one goes down a military and/or spycraft path, the dynamics are somewhat different, particularly with regard to absolute protection of human life and commercial viability, but the dollar figures are as large or larger, and competing priorities for those dollars provides enough of a commercial-like environment that the end result roughly parallels commerciality. And, for those things that are truly breakthrough, but kept secret, if any, they won't see the light of day for a long time in anything that is truly general, like technology-only based AGI.)

There are actually several examples of this commercial disincentive impact in history; some are described in "The Innovator's Dilemma", which you've read, I believe. The corollary is also laid out in that book, whereby disruptive technologies obliterate established technologies; but, I would argue that those fall within a narrowly defined area. It is exactly this narrowness that makes the extrapolation of disruptiveness to AGI and ASI fallacious.

Artificial Narrow Intelligence (ANI) exists today. I, personally, don't like this label since it conveys a notion of more than what is really happening. In the retail consumer world, the term "smart" is applied (to phones, cars, homes, credit cards, etc), which is just as misleading. All of these things rely on computer programs, which are nothing more than an accumulation of rules that have been codified by someone(s) based on what they can imagine and then describe in mind-numbing detail. Every one of those rules, when dissected to its most basic level, is a series of absolutely binary branches in a logic tree. Despite all of the advances in the components of computer technology over the years since ENIAC was introduced in 1946, 70 years ago, it still comes down to ones and zeroes (true/false, on/off, yes/no, black/white). There is no 1/2 or 2/3 or 17/59 or whatever in a computer; there is no "hmmm", there is no "what if", there is no "sorta kinda". There is no thinking, wondering, trying, and imagining ... there is no intelligence. They are not smart; in fact, they are incredibly stupid. A person can be considered insane if he does the same thing over and over, somehow expecting the result to change. A computer, by contrast, can't be considered insane because it will do the same thing over and over and over and over without any expectation ... it just does what it's ones and zeroes dictate ... nothing more, nothing less. To date, all that has happened is to make a whole lot more, and much faster, ones and zeroes available at such low cost that everybody can have some. Until such time as computers can really do a "what if" on their own, we will be constrained by what humans can imagine and describe with enough clarity to put it into ones and zeroes. If we did want a computer to take control of the world's resources and eliminate humanity in the process, we'd have to figure out all the steps needed to do that, program it into a computer (or lots and lots of them), convince all of humanity to not do anything to try to stop it, and then let it try. The likely result would be the computer crashing due to some physical failure, or the program either halting or entering an infinite loop, because we humans would have overlooked something in the process and failed to account for it in the instructions we gave to the computer; so, it would do exactly, and only, as it was instructed. (I haven't included any discussion regarding "learning" computers, like IBM's Watson, since it doesn't really alter the underlying principles above ... humans still have to instruct the computer on how to do "pseudo" learning.)

Granted, we do have an increasing variety of sensors (e.g. lasers, accelerometers, microphones, photo-optics, spectrometers, etc) affordably available to us to translate environmental into ones and zeroes on our behalf; but, at the end of the day, it's more of the same stuff. So, as things stand, cheap computers enable our tools to be more effective; but, they are just tools. (Appropriately marketed, some people might buy smart hammers; some did buy pet rocks!)

The article mentioned advances in nanotechnology and the possibility of nano-robots doing material manipulation at the atomic level. Not sure I can say much about this other than to point out that not a lot has changed in quite a while when it comes to elements on the periodic table and/or the interaction between them. And, while we've gotten pretty proficient at fission reactions (e.g. nuclear power plants and bombs), we still haven't cracked the code for any sort of practical fusion reaction, which is where some of the theoretically interesting things lie. More importantly, I think, is what's happening at the sub-atomic level; which, by the way,

is far outside of what is even being postulated for nano-bots. Quarks, muons, leptons, “God particles” aka the Higgs boson, etc are very new observances in the grand scheme of things. But, we can only infer their existence by measuring distortions in a micro-universe and comparing the collected data against theoretical models. (This is another one of those things that costs billions and takes years to make even the tiniest of inroads.) And, even if we could grab hold of one of these sub-atomic things, we couldn’t measure it because the act of touching it changes it, at least according to the theories. On the “God particle” front, the exciting news is that those chasing it are pretty convinced that they have proven its existence. Their predictions were that it would have one of two behavioral patterns, each of which pointed to a competing, yet well-defined theory of the deep structure of the universe; however, the data indicates that its behavior is right in the middle, so it might point to a completely different fundamental structure, or it might mean that both exist, somehow. (Wow, maybe they’ve found something that isn’t black or white, which could be interesting as a basic building block for a truly thinking computer!!! Okay, I just jumped light years past Mr. Spock, so not by the year 2060.)

Biological computing is being dabbled in. The concept is fascinating as it leverages the self-replicability and adaptability of cells (think stem cells) to grow, repair and “learn”. I use the term learn because what happens in our brain as it develops/grows is cellular at its basic level, and it encompasses learning and thinking, not to mention emotions and the five (or six) senses, and other pretty phenomenal capabilities ... in other words, intelligence. Cells are clearly not limited to zeroes and ones, they can and do represent multiple states, particularly as they relate to other cells and stimulants (e.g. hormones, enzymes, electrical fields, light, temperature, oxygen, etc). If one could harness all of this, a truly intelligent computer could be envisioned ... today we tend to refer to this as a brain.

I would argue that what we are doing today is essentially hybrid computing ... that is to say, along with our brains we use computing technology to help it manage and solve complex problems more efficiently. However, the interface between our brain and the computer is incredibly primitive (we type on keyboards, or speak a limited vocabulary, or a limited dialect); though there are promising developments in the area of what I’d call bionics, where “smart” limbs are responding to bio-triggers to effect action. Controlling the functioning of a bionic arm by simply thinking about the desired actions is a great example. Over time it is reasonable to expect that cognitive interfaces to computational devices will achieve practical levels. This is where I think we have the best chance for getting anywhere close to AGI. Voice synthesis and recognition technologies are roughly 40 years old, and they aren’t that sophisticated, so cognitive control and feedback mechanisms at a practical level (don’t forget the feedback part since that’s the only way the results of the computer’s work can be cycled through the brain for intelligent thinking, what-iffing, and the issuance of further commands to the computer) are a long way off. Not 2040 or 2060, and probably not before 21xx. To be clear, I am not suggesting a disembodied brain a la sci-fi movies; but rather, a progressive evolution of the man-machine interface continuum to the level where the elegance of the interactions is such that it feels natural (a lot of that, by the way, simply means that it has become habitual). The difference between typing on a keyboard and “asking Siri for help” is an excellent example of this

progression. I would argue that a bi-directional audio (speak/listen) interface between man and machine falls short of AGI, and certainly doesn't qualify for ASI.

Viruses have been mutating for hundreds of millions of years, perhaps one of the best, truly exponential examples in the universe; their ability to try literally billions of combinations simultaneously is fueled by the absolute number of viral entities on the planet, the diversity of their host populations, the interactions between these hosts (e.g. mosquitos, humans, birds, pigs, ticks, deer, etc, etc) and human facilitation with such things as air travel, antibiotics, our penchant for exchanging bodily fluids, and, ultimately, their total focus on their own survival without any concern for any other being on the planet. They are, in many ways, the embodiment of the uncontrolled computer in the article. And just to make things more interesting, we cook up more esoteric variations in the lab with genetic engineering that may not have ever occurred naturally; which we end up introducing into the wild, either intentionally ("controlled" human trials) or unintentionally ("oops, we lost a couple of vials"). There's no question that we do not now, nor will we in any expected number of lifetimes, understand this homegrown alien world to the level needed to effectively manage/control it.

Meanwhile, the most basic behaviors in nature continue, as they must. That is: most animate beings are genetically programmed to eat, sleep and procreate, while viruses are programmed to eat, mutate and replicate. By the way, this might be the best use of the word "programmed", as it is really the original description of natural behavior ... computer lingo merely derives from our attempts to imitate life.

So, given the negative human impact we've seen from viruses in the past, it isn't hard to imagine that a particular series of mutations could wreak havoc upon us, even to the point of extinction. On the positive side, though, in this age of biotechnology, genetic engineering, leveraging what we see/find in the laboratory called nature, and pure discovery, we are fairly well positioned to detect and respond to biological threats as they arise, via quarantines, sanitation, preventatives and curatives. Further mitigating against the odds of viral-based apocalypse is the fact that viruses hold no malice, have no intent, and can't learn to leverage trends ... what they do is totally random; we, on the other hand, generally act with learned intent, which gives us a slightly better than random chance of success.

At the end of the day, it isn't about biology or physics or computer science; these are merely abstract terms we utilize to describe the vantage point from which we are viewing a slice of the universe. Once we get down to the subatomic level (i.e. the things that make up electrons and neutrons and positrons), there is no distinction. Viruses are doing things at this level (we call it molecular biology, though we should probably call it subatomic biology) and they are doing it randomly ... so, they don't know what they are doing. We, meanwhile, are knowingly messing around at the subatomic level, but we don't really know what we are doing either. Should these two "idiotic" efforts unwittingly collide in a bad way, all bets are off. On the other hand, should these two "idiot savants" stumble together into something wonderful, then there could be a real "aha" moment.

Why the subject line: Artificial Intelligence and Grandkids???

Let's ask them the same questions posed in the article in 25 years, when they're out of school, and it's 2040, and they've moved on to something more elegant than Siri, and they don't even remember what a keyboard is (unless it's in black and white and has octaves) ! Maybe they'll have an answer. Some will undoubtedly declare that AGI has been achieved. Maybe so, but I doubt that it will have reached the level described in the definition provided. I'd say that 2140 is a more reasonable target timeline. As for ASI, that's a whole different ball game ... I couldn't begin to hazard a guess for when, if ever.

That's my two cents, adjusted for inflation ...

Anon