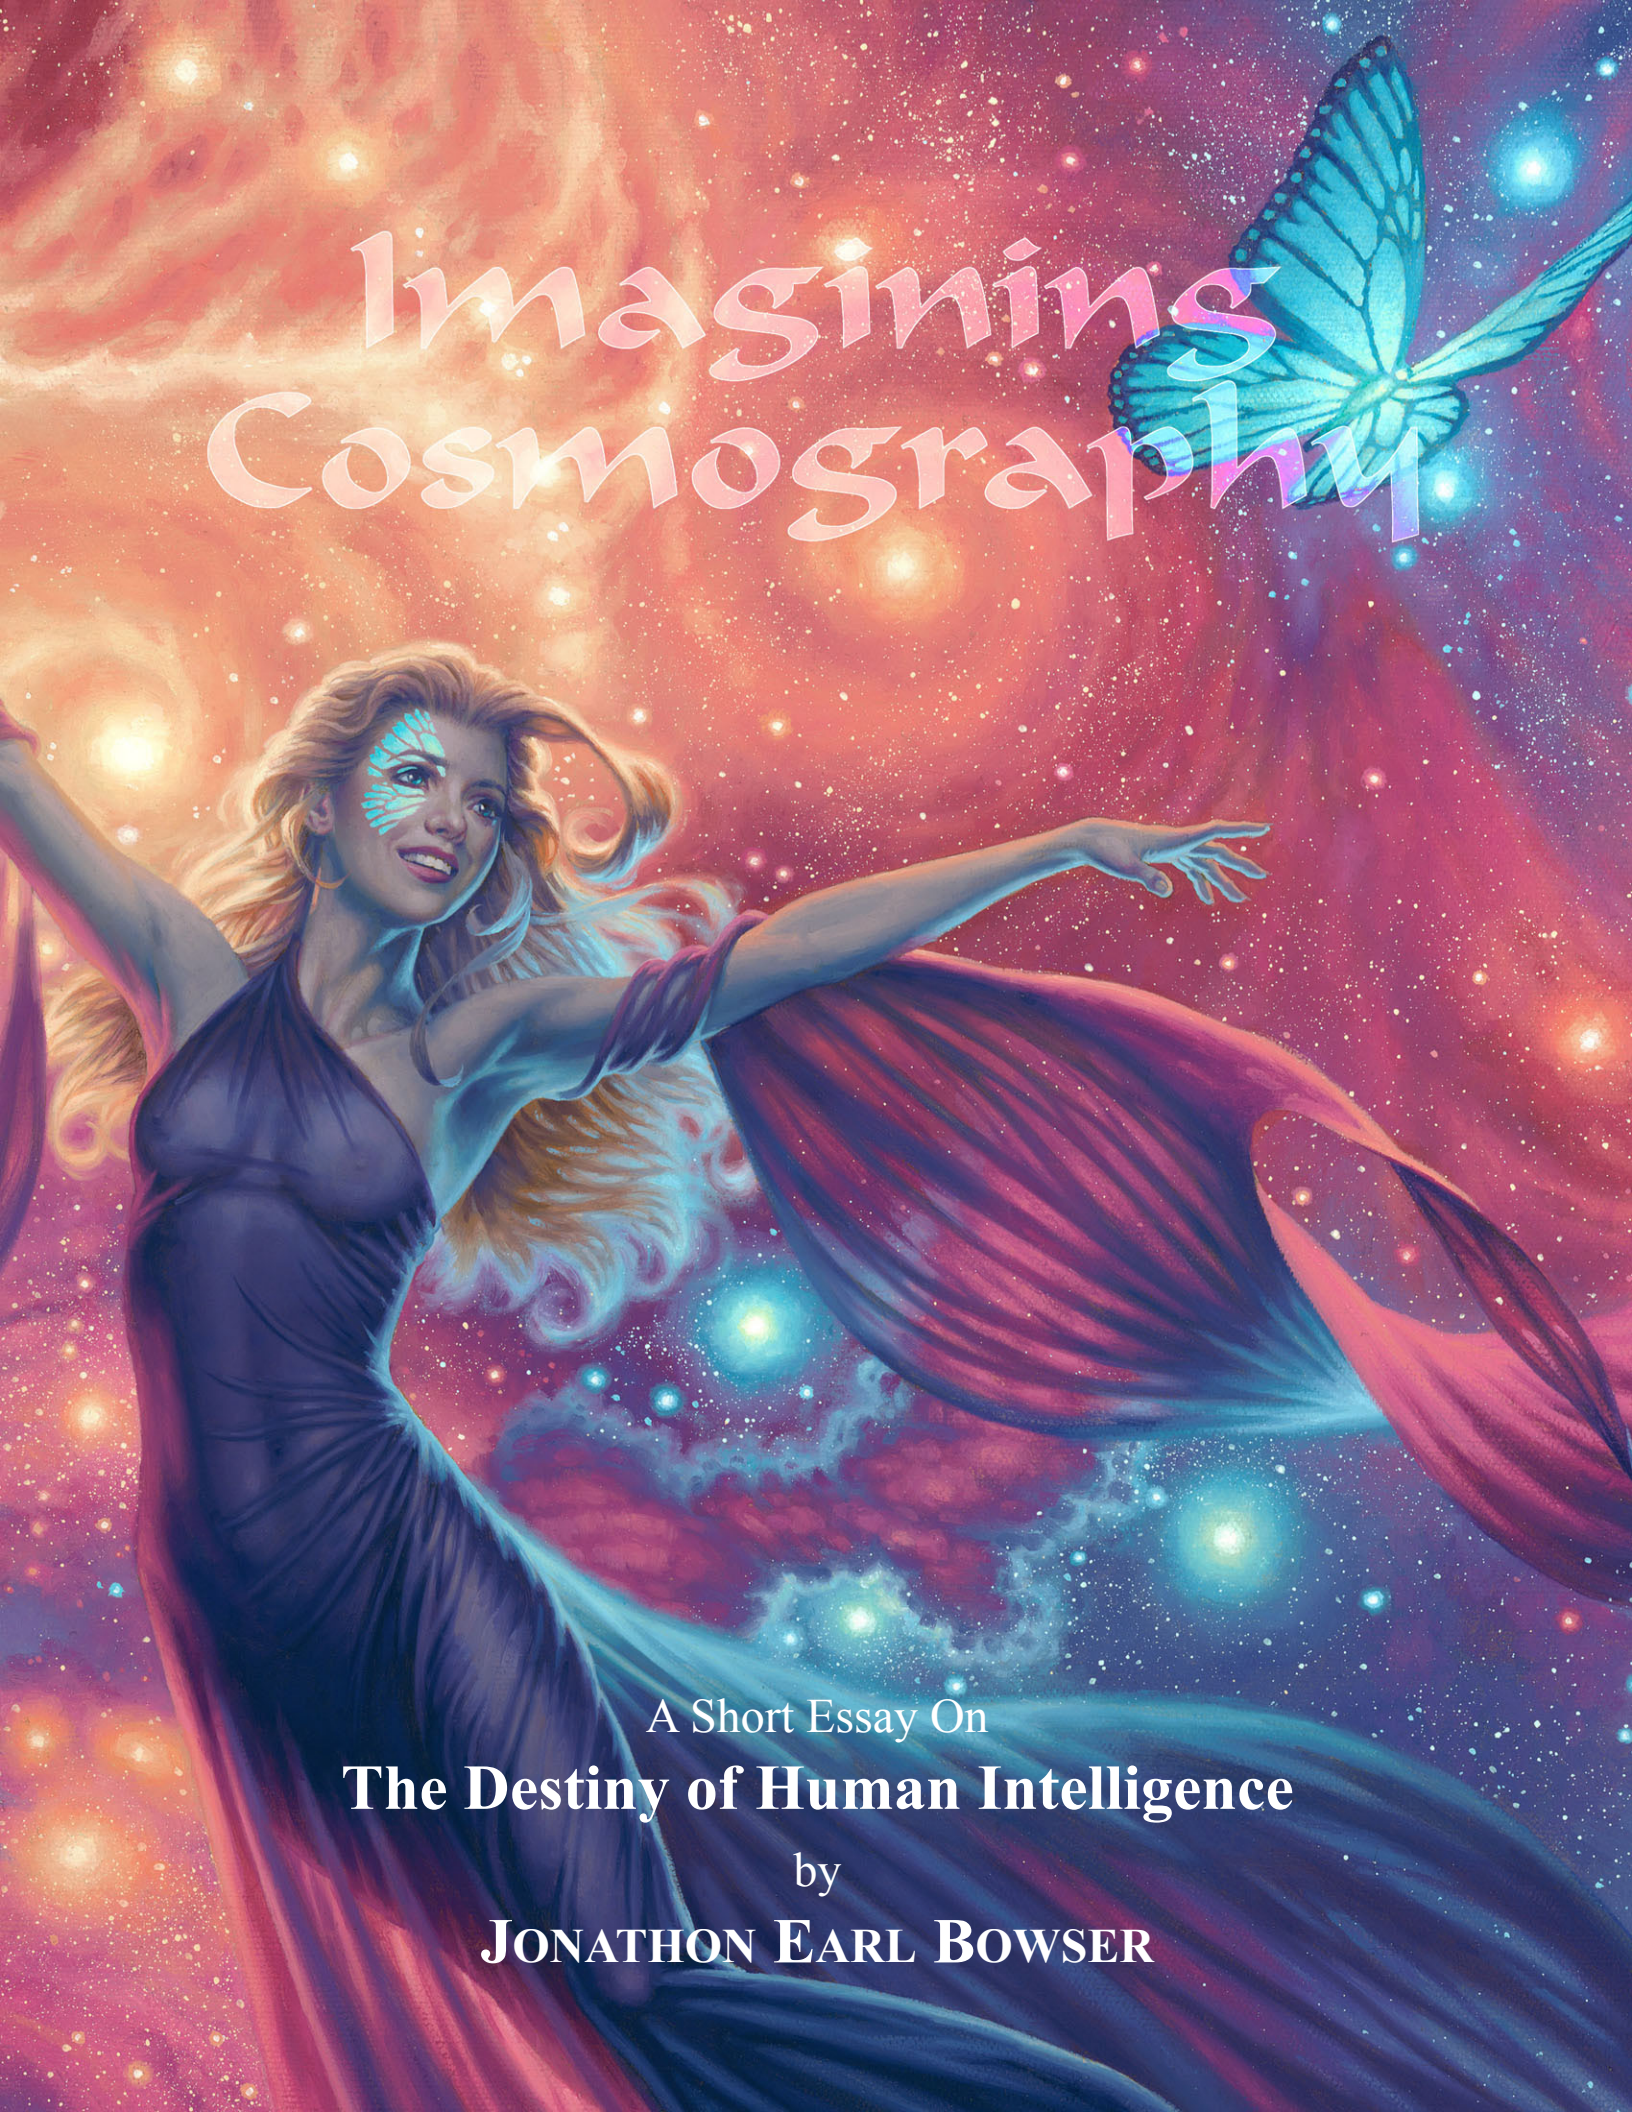

A Short Essay On

The Destiny of Human Intelligence

by

JONATHON EARL BOWSER

Kurzweil's Optimism

...we can expand out into the rest of the universe… sending…nanotechnology infused with artificial intelligence…to expand the overall intelligence of our human machine civilization. The universe will wake up; it will become intelligent.

Ray Kurzweil

When I first heard Kurzweil's inspiring vision of the future, as an artist I wondered, *what would that look like*? Not literally, a picture of billions of tiny robots drifting into blackness, but poetically, an image to suggest the Second Human Diaspora as we become a great interstellar exodus of creative intention in the cosmos. Before discussing some of those ideas near the end, we should begin by trying to answer an obvious question about this audacious prediction. Is Ray Kurzweil - inventor, futurist, author, all-around genius, and a director with the Google Brain Project – correct in his assumptions? Is the destiny of human intelligence in the stars? The almost spiritual undertone of his faith in technology might make it easy to dismiss him as some kind of evangelical huckster. But his erudition and accomplishments (19 patents, a Presidential Medal of Technology, and much more) demand that we take his predictions seriously.

Kurzweil's projections are based on at least two assumptions that may indeed be reasonable, but it is important to recognize that considerable speculation is involved. The first assumption is that *Moore's Law* will hold for the foreseeable future. Gordon Moore, founder of *Intel Corporation*, observed in 1965 that processor speed and capacity had been doubling every 18 months or so since the IC chip had been invented at *Texas Instruments* in 1958. This trend in microprocessor development has continued uninterrupted since then, and human-level capacity could be achieved by 2045 – again, if current trends continue. Enormous financial and human resources are dedicated to this task, and technological advancements happen rapidly. Past success is no guarantee for the future, but the prospect of trillions in potential wealth has a very powerful way of focusing a civilization's attention.

The second assumption is a little harder to nail down. It seems quite likely that we will one day have a computer-brain fast enough to run a neural-like cluster of processors with a complexity rivaling the human brain; what is less certain is that we will have a computer-mind able to use it. In 2016, nobody really knows what a mind is, which will make building one problematic. The matter of consciousness is so mysterious that eminent MIT linguist Noam Chomsky once said, "We don't even have a theory about what a Theory of Consciousness would look like." That sounds like an insurmountable obstacle, and perhaps it is, but there is still good reason to be optimistic. Human brains, and the minds that run them, get built by nature all the time. The world is filled with them, so we know (or strongly suspect) there's nothing magical or contrary to the Laws of Physics about constructing a human-level mind. If we study the operational prototypes (us) closely enough, image the function of the brain in ever finer detail, carefully document how natural processes assemble it piece by cellular piece, eventually we'll be able to retro-engineer a satisfactory approximation. Or so the prevailing wisdom believes. This kind of very advanced science doesn't sound immanent, and C3PO by 2045 is perhaps overly optimistic, but it does sound *possible* – eventually.

The promise is indeed an alluring one: download your consciousness into a small interstellar probe and explore the universe forever. A significant percentage of humanity will find that promise irresistible - certainly preferable to the dead-forever alternative - and the obvious trend of technological development is toward increasingly human-like computers. Well-organized research teams of the smartest people in the world are working with almost limitless resources, motivated and inspired by the challenges of a hyper-competitive intellectual environment. Surely, with arbitrary vistas of time to achieve our ambitions, success in this god-building endeavor is a foregone conclusion? Surely, as each algorithmic improvement accelerates the pace of subsequent development, a critical threshold must eventually be crossed - in 30 years or 300? Someday, everyone will become the captain of their own personal, self-replicating, humanity-dispersing starship! Right?

Perhaps. But, as technology evolves apace, deeper problems in these not-quite-so simple assumptions emerge…

Darwin's Exigency

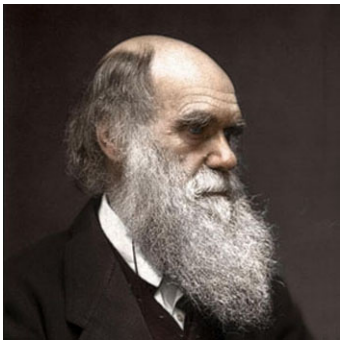

It is not the strongest of the species that survive, nor the most intelligent, but the one most responsive to change.

Charles Darwin

Life, apparently, is not the only thing that evolves. Technology of the kind that meets everyone's expectation of the word - electrically-driven machine-made tools with microscopic precision - has only existed for about a century or so. When compared to the 5000 years that have passed since the dawn of the historical period, we've moved forward in the span of just one very long human lifetime from clunky steam engines to robot-laboratories doing molecular experiments on the surface of alien worlds. And still, our most formidable institutions of education and research are relentlessly advancing our understanding of the materials and processes of nature and how we might harness the limits of what the universe will allow us to do for new, even yet-to-be imagined, purposes.

Developments in nano-machine, biological and genetic, cybernetic, micro-processor, and programming technologies will soon present humanity with powers that, even in the 21st century, sound more like magic than science. If you should live just another 30 years or so, here are a few of the miracles you can reasonably expect to see, technologies that are *already* under development:

It is quite likely that you will be able to purchase *nano-augmentation* (billions of atomic-sized robots swarming purposefully through your body) specifically tailored to your unique genome that will significantly improve your strength and durability, endurance and performance, resilience to disease and toxification, and even longevity. We will live longer, healthier, more vigorous lives.

It is quite likely that if you should lose a limb to accident or an organ to disease, a synthetic replacement indistinguishable from the original – machine-made or genetically grown – will be available. Some may even choose to upgrade otherwise satisfactory native parts for superior artificial replacements. We will eradicate disabilities and vastly surpass the limits of what humans can presently do.

It is quite likely that computer processing power and speed will be able to place the entire corpus of human knowledge into an efficiently usable, portably wearable device. It remains to be seen whether such machines might ever be able to purposefully think for themselves, but it is inevitable that personal data-chips will make more knowledge than any human has ever possessed instantly and *usefully* available to our brains. Revolutionary syntheses of radically different forms of thought will become common, and world-changing. We will all be *much* smarter than Einstein.

Many people will regard such physical and cognitive enhancements as contrary to the natural order; they may see any kind of anthro-engineering as sacreligious and will want no part of it. Many more people will not be able to afford such upgrades. And so a chasm might open between humans and what some are calling *trans-humans*, a gulf that one can easily imagine will rapidly fill with envy and suspicion, and violence may result. But it is clear that any new class of super-intelligent humans, whether large or small, will accelerate such divisions; within some very small period of time (in evolutionary terms) unaugmented humans might seem in comparison as intellectually limited as chimpanzees do to us. And these meta-humans, capable of feats of logic and computation incomprehensible to retro-humans, will rule the world (perhaps benevolently) because they will surpass us in every field of human endeavor. Chimps are not capable of governing humans, who will not be capable of governing trans-humans.

The truly astounding rewards promised by such technology pull us down this path irresistibly, and there will be no turning back from such awesome prospect. Perhaps, like many luminaries in the field of Artificial Intelligence, you also believe there should be some serious debate about this project, a deep consideration of the wisdom – or lack thereof – in pursuing such power. If something should go terribly wrong, if our synthetic savants should in any way deviate from the exquisitely exact purpose we intend for them, what hope would we have against hostile atomic armies under the control of a super-intelligent AI commander? But no such bird's-eye circumspection is likely, it can be argued, because an inexorable evolutionary calculation is always busily grinding away just below the surface of polite conversation: *maybe I don't really need that power, but I can't let my competitors or enemies get it first*!

This is a strategic imperative that applies to everything from bacteria to baboons, but is easiest to see

in the actions of intentional plan-makers like us: if you recognize a beneficial resource or any other potential reward in the environment, you must use it for your advantage before it is used against you by another for their advantage. Whether in the competition for limited resources or the struggle for life and death, it is the player with the most advantages that prospers. There are many successful co-operative strategies also, but as the stakes of the game get ever higher, co-operation becomes increasingly risky and difficult. I have no doubt that co-operation will define the early phase of this research. But, as we shall soon discuss, the realization of strong AI will only happen once and will be the most significant moment in human history. The cost and profound security implications of advanced AI will require the participation and supervision of national governments; their record in this regard is significantly less than adequate, and co-operation in the late stages of research may be impossible. And so the entirely natural processes of natural selection and survival of the fittest continue into the digital age…and compel us at breakneck speed into unexplored realms of god-like power.

Jung's Conundrum

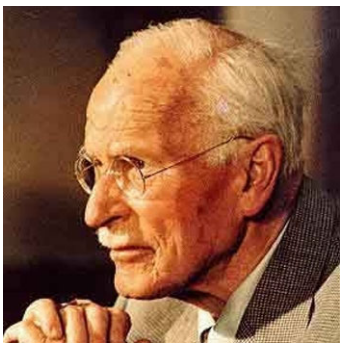

One does not become enlightened by imagining figures of light, but by making the darkness conscious.

Carl Jung

I hope you will pardon a brief digression here to ponder an aspect of mind that is rarely, if ever, discussed, and I think it's a question of some importance if we mean to actually build these things one day in the rapidly approaching future.

In mythology there is a very common theme that involves an intrepid adventurer entering into a place of mystery: a cave, a dungeon, a forbidding forest, the bottom of the sea, the belly of a whale, or any dark and mystic enclosure where the reassuring light of the everyday world cannot reach. This ancient and popular motif in sacred stories from around the world is symbolic of a return to the womb. Now this is not a return to the womb of the body, which sounds unpleasant for everyone involved, but a return to the *womb of the mind*, that seminal bud of potential that is the pre-conscious state - who we were *before* we became somebody.

This imaginary voyage to the empty darkness which prevails at the genesis of consciousness seeks to visualize the new and naked psyche as an empty room into which we can then pile our acquisitions of experience. The original structure of this room is an inheritance, a gift from our parents over which we have no control. And the early shape of this mind-bud will be receptive or conducive to certain contents but not to others. It will accommodate some contents easily; others may obviously not fit at all, but we are obliged to begin the lifelong process of filling this room long before we are wise enough to know what contents might best suit a room such as the one we were given before birth.

To a certain extent we can exercise some choice in what goes into the room - this incarnation of our psychic identity - but very soon it is the other contents of the room that influence our choices for subsequent acquisitions. Behavior reinforces behavior, and repetition invites repetition. If early choices have been made unwisely, and they often are, then the character of the room continues to evolve in a misguided direction ill-suited to the true nature of the room itself.

In mythology, the hero's quest is to render invisible the chaos of the conscious contents of the psyche, all mere accidents of our accidental experiences, and boldly step into the not-yet-illuminated darkness of not-yet-imagined possibilities. It is a process of, not accepting what was given, but finding that which might be taken with deliberate intention, seeing the important empty niches in the timeless room of the mind that can be artfully filled with a precision-crafted treasure.

Now that fanciful tale doesn't sound like it has very much to do with a discussion of Artificial Intelligence and the Future of Humanity, but I think it does have some relevance to the challenges of human cognition. And if we're going to build these synthetic super-brains with super-minds, we'll have to figure out how to build, not just any old functional mind we can cobble together, but a mind that is perfectly tuned to the architecture we give it. It is quite likely that any such architecture will be understood imperfectly, designed in large part by opaque "black-box" AI systems that understand the big-picture concerns even less than their human masters. Anatomically and physiologically the brains of Einstein and Stalin were identical; it will require a profound discovery indeed to deeply

understand the ethereal differences of mind that so importantly distinguish them.

The closer one looks at the challenge, the more one considers the risks of failure, the more intractable the whole endeavor seems…

Bostrum's Pessimism

Far from being the smartest possible biological species, we are probably better thought of as the stupidest possible biological species capable of starting a technological civilization—a niche we filled because we got there first, not because we are in any sense optimally adapted to it.

Nick Bostrum

Professor Bostrum speaks extemporaneously in long, syntactically elaborate sentences, about deep and abstract subjects, with a very quick and Swedish-accented voice. He is a philosopher (with a background in mathematics, physics, and computational neuroscience) and the founding director of the Future of Humanity Institute at Oxford, where he and his brilliant colleagues contemplate the existential risks of new technologies, and strategies to mitigate these risks.

Current human civilization, Bostrum observes, is anomalous: historically, geologically, cosmologically. For most of our 100,000-year existence as a species, humans lived short and brutish lives in squalor, near starvation; only in the last 100 years has any significant percentage of humanity enjoyed comfort, leisure, and security. Much of earth's past is characterized by extensive volcanic activity, hemispheric glaciations, atmospheric upheavals, and mass extinctions; this current agriculturally fertile period of geologic quiescence was preceded by a long epoch of frigid climate utterly incompatible with civilization. And the universe is overwhelmingly vacuum and radiation; the gossamer-bubble veneer of habitable biosphere on this rocky planetary crumb is an incalculably small speck of the entirely unsuitable cosmic real estate available.

To be living here and now, safe within the protective cocoon of technological society and enjoying the agreeable weather of a warm watery world, is *highly* anomalous; as time goes forward, it becomes ever less likely such improbably good fortune will endure. The current unstable state must eventually settle into one of only two possible stable states. The first option for a stable human condition is, as proven by the fate of 99.9 % of all species that have ever lived on earth, extinction. Given enough time, a global catastrophe - natural or man-made - is certain. The second option, according to Bostrum, is to inherit our "cosmic endowment," a period of cosmic colonization where the intelligent machine-progeny of human civilization disperse into the deep universe for billions of years. From an existential perspective, a colonizer state is a stable one because, with universal distribution, humanity becomes immune to even galactic catastrophe. It's rather like the crumbling pinnacle upon which our fate is balanced is dissolving below our feet; at some perhaps not so distant point humanity will be forced to leap in one direction or the other. Professor Bostrum, like almost everyone, believes the latter option is infinitely preferable to the former.

Bostrum imagines technological evolution as someone pulling colored balls from an urn. Most balls are white, indicating a beneficial discovery – like mathematics or electricity. Grey balls indicate dangerous technologies that may seem of questionable value – like torture devices or nuclear weapons – things we might have preferred not to invent but have not (yet) killed us off. A black ball would represent a technology that always drives its discoverers to extinction. We haven't plucked such a ball yet, but we don't know that there isn't a death-ball in the urn. Nuclear weapons turned out to be very difficult and very expensive to build; so far that developmental barrier has kept them out of the hands of lone psychopaths with malicious intent. But if we should one day discover an *easy* way to acquire such city-destroying power, a method of destruction so simple that anyone who wanted it could make it quickly and cheaply, well, that would be the end of cities and civilization.

And technological civilization might be a one-shot opportunity because the first tech-izens (like us) will probably exhaust all the easy resources in a first industrial revolution. That is, all the oil, gas, coal, minerals, metals, etc., that you can get at *without* advanced technology – the stuff near the surface – is gone forever. How would a resurgent civilization recovering from the selection of a black-ball technology ever find the deep resources (which might remain) without the kind of advanced technology that's made only with

materials you can no longer find. And so a second technological civilization eventually rising from the ashes of the first (which could have destroyed itself in any manner of different ways) seems highly improbable. Human inclinations to conservative feudalism and self-limiting superstition impeded the advent of our industrial revolution for at least 1500 years; a permanent inability to gather the energy and materials that technological development requires would probably delay it forever.

But let us optimistically assume there are no 8-balls in the urn, and that dangerous new technologies will be carefully developed and contained by large teams of responsible, well-vetted and altruistic people. It is generally thought that once we have developed an artificial intelligence capable of doing any human-level task competently, then it is quite likely that such an AI would continue to use its AI-building prowess to improve itself at a digitally-accelerating rate. It would quickly become an advanced *superintelligence*, radically superior in every way to anything humans can imagine. It is also quite likely that this event will happen only once. Competitors in the race for AI may be months behind the team that gets there first and when this *intelligence explosion* occurs - in seconds or weeks - superintelligence will be achieved long before any competing system can also reach the AI threshold. The race for AI is Winner-Takes-All. With no vaguely comparable intelligence anywhere in the world, this new superintelligence will determine the nature of all subsequent technological development; the *Supreme Inventing Machine* will be the last invention humans ever make.

Perhaps the most vexing challenge in this matter is something called the *control problem*. If a team of mice wanted to contain human agency in the world, it is easy to imagine just how comically inadequate their best efforts would be. It will be far too late for us to contain the AI after it achieves superintelligence, so it is critically important to set up initial conditions that will optimize the prospects for a beneficial outcome. According to Bostrum, recent progress on this question has resulted only in, "a deepening appreciation for just how profoundly difficult this problem is." Even with a (still-imaginary) *control solution* in place, are appropriate safety precautions even possible in such a race? When frantic reports from our military-industrial espionage come in, telling us that our competitors are very near to a breakthrough in machine intelligence, will we still delay our program to wait for the results of research into the mushy-headed question of whether our AI has hidden aspects of StalinMind? We might hope for some co-ordination between America and Europe in this endeavor. We should be less confident about the prospects for caution and restraint in other parts of the world. And our history in this regard does not always inspire optimism.

During the height of the cold war in the early 60s, the Russians and the Americans each had about 30,000 nuclear weapons, with an average yield for each of about 10 megatons. That's 600,000 megatons total, which is 600 billion tons, which is more than a million-billion (1,000,000,000,000,000) pounds of explosive potential - a half-million pounds of TNT-equivalent under the feet of every human on earth at the time - all aimed at the civilian populations of the world. The nuclear weapon was invented in 1945, and within 20 years the politicians of just 2 countries had 2 railway cars of dynamite strapped to every man alive…a 6-billion-car train of death wrapped around the world 2000 times like a global suicide vest. It's difficult to comprehend the monumental perversity of our collective dementia, but perhaps we should just be grateful for the fact our deranged executioners haven't yet decided to annihilate the planet. Hallelujah.

It seems we are in a desperate race with our own monstrous depravity: one team seeks to invent the next even more terrifying super-weapon that might quickly escape our flimsy control; the other seeks to invent the machine-assisted wisdom to help us see at last the colossal magnitude and folly of our misallocation of resources. Odds-makers are offering less than even-money…

The same evolutionary strategies that allowed humanity to survive the vicissitudes of brutal nature through unknown millennia of pre-history now incline us to immanent self-destruction. We need a new human nature and we need it quickly. Optimists believe we can learn *Cooperative Intelligence* eventually; pessimists fear we can't unlearn *Competitive Stupidity* fast enough. Carl Sagan's famous question to our first space-travelling visitors (spoken by Jodie Foster in the film *Contact*) finds the bitter heart of the matter: *How did you evolve as far as you have and not destroy yourselves*?

Darwinian forces of natural selection must obtain in any biological ecosystem anywhere; without rigorous selective pressures, without a legitimate reckoning of fitness, without compelling environmental reasons to adapt, without a merciless resolve

to survive, evolutionary progress toward anatomical and neurological sophistication just wouldn't occur. It seems certain that any intelligence that evolved elsewhere in the cosmos must have, at some time in their developmental trajectory from murderously competitive to creatively cooperative beings, had to run this same harrowing gauntlet of technological self-destruction. It would indeed be edifying to learn just how existentially treacherous this path is. It's not unreasonable to wonder: Has anyone anywhere passed safely through it?

Hoyle's Improbability

Once we see…that the probability of life originating at random is so utterly miniscule as to make it absurd, it becomes sensible to think that the favorable properties of physics on which life depends are in every respect deliberate...

Fred Hoyle

Sir Fred Hoyle was an astronomer famous for his Nobel-winning theory of stellar nucleosynthesis, which correctly explained how heavy elements are cooked up in the explosive deaths of stars; he was notorious, too, for his skeptically derisive name for the explosive origin of the universe: the Big Bang (which, it seems, got him kicked off the team that actually received Nobel's loot). But he also did an interesting calculation about the origin of life that I'm going to call *Hoyle's Improbability*.

Once you have a complex molecule capable of making copies of itself by purely chemical means, the process of life making more life is comparatively easy to understand. But it remains a deep mystery how you get that *very first* self-copying molecule. We don't see simple atoms autonomously assembling themselves into significantly more complex, never-mind self-replicating structures, and Hoyle reasonably wondered about the likelihood that the requisite atoms would *just happen* to find themselves in the right place at the right time. Without the proper organic chemistry to put those atoms in the right place, *how did they first get placed*?

It seems that the smallest replicator-molecule we can imagine, one having enough parts to allow basic chemical shuffling that could also involve accidental reiteration, is a simple strand of RNA – a precursor to DNA that consists of about 20,000 atoms arranged in a precise 3-dimensional matrix. Hoyle's calculation tallied up the parts - bits of hydrogen, nitrogen, carbon, oxygen, and so on – and their positions, to determine the odds that this proto-replicator could assemble by sheer chance alone. He came up with an *exceedingly* unlikely figure. The chance that those 20,000 atoms would randomly collide, into a perfect bird's nest cluster of sufficient complexity to initiate the self-sustaining auto-replicating chemical reaction we call Life, is 1 in $10^{40,000}$.

Now this is a quite extraordinarily large number. How big is it? Allow me to present a thought-experiment to illustrate its *stupefying* bigness. This experiment could not actually be done, of course; it is merely a way of starting to think about a number beyond comprehension. There are, we believe, about 10^{80} atoms in the observable universe, and physics says that the smallest possible unit of time – *Planck Time* – is 5.39 x 10^{-44} seconds…which I'm going to simplify to 10^{48} possible cosmic ticks per second, just to keep our arithmetic easy.

Now let us imagine a colossal *atom-smasher of the imagination* where we collide every atom in the universe together 10^{48} times every second, and with each cosmic scrunch we generate a possibility that some 20,000 atoms somewhere in the experiment has happened to get randomly assembled into the desired RNA-replicator fashion. Now, 10^{80} atoms shuffling 10^{48} times a second, is 10^{128} possible atomic combinations tried every second, or about 10^{135} experiments per year. So, how long would we have to run this experiment before we had rolled the atomic dice $10^{40,000}$ times, before we had a reasonable even-odds prospect of achieving the specific orientation of specific atoms we seek? Well, if you took every atom in the universe and smashed them together a trillion trillion trillion trillion (10^{48}) times a second, and you did that for a trillion trillion trillion trillion years, you'd have run the test…just 10^{183} times.

I've taken you on a bit of a wild goose chase, it seems. If the odds of life forming by chance are 1 in $10^{40,000}$, that means there's no chance at all…and yet here we are. Life obviously *didn't* fall out of an impossibly fast and impossibly long experiment of the imagination; life happened almost immediately, perhaps less a billion years after the formation of the earth, based

upon nothing more than simple chemical laws that favored a few possible combinations over a near-infinite number of other impossible combinations.

Why is that interesting? Because there is no prospect for life happening on a young world by chance alone. None. Life happened because *the laws of chemistry compelled it*. And what does that mean? It means that life, along with the universal laws of chemistry, *must be everywhere*. It is overwhelmingly likely that every stable star system with a small rocky planet in the habitable zone where liquid water can exist, will have at least simple life on it. Multiple stars may be more complicated, but if just one star system in a hundred meets this easy criterion, there should be *a billion living worlds* in this galaxy alone. We haven't heard from any of these living worlds yet, but it's just a matter of time. They've gotta be everywhere! Right?

Fermi's Paradox

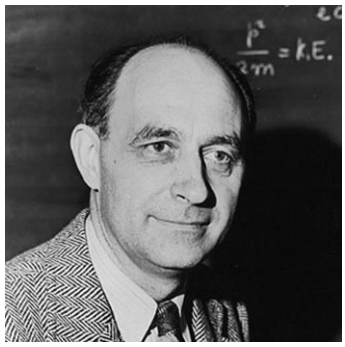

Where are they?

Enrico Fermi

The universe is large – about 600 billion-trillion miles across. Given that just our Milky Way galaxy could have as many as 100 billion habitable worlds (one habitable zone per star), and there are perhaps 100 billion such galaxies in the cosmos, then we should reasonably expect that enough rolls of the intelligent-life dice have occurred and there should be other technological civilizations out there. So where are they?

The universe is old. Given that the first second-generation stars (with their planetary rings of life-enabling heavy elements made in the collapse of first-generation stars) were coming into being almost 10 billion years ago, there has been plenty of time for intelligent life to arise, evolve to levels vastly beyond human, and traverse the cosmos even at 10%-luminal chemical-rocket speeds. So where are they?

The universe is entirely natural in appearance. Given that there is an obvious and easily discernable distinction between wild and civilized places on earth and that we should expect at least *some* of the other intelligences out there to have comparable inclinations to large-scale resource extraction and energy management, then when we look to the sky with our sensitive measurement machines we should reasonably expect to see – even across intergalactic distances - the characteristic electro-magnetic signature of stellar, interstellar, or even galactic engineering akin to the geo-engineering that humans display so abundantly here on earth. So where are they?

And, once again, we are faced with the Darwinian Exigency of the Prisoner's Dilemma: maybe very nearly all advanced civilizations decide that there's no need to pursue a program of galactic engineering and colonization – that's just something that the primitively violent simians on Sol-3 believe. But there are billions of competitors and only *one* of them has to be the tiniest imaginable amount paranoid about the unknown intentions of another mysterious planet…for them to be *compelled* to colonize the universe. This is Darwin's Law: eat or be eaten, colonize or be colonized. No one *wants* to colonize (our sentimental hearts suppose), but everyone *must* colonize or face subjection by other colonizers. Straight-forward calculations demonstrate that an initial launch of self-replicating probes of a kind already possible with current human technology could occupy the observable universe in just a few million years, and billions of planets had billions of years to complete their exploration and distribute their culture throughout the universe before the earth even existed. So where are they?

Legend has it that, when discussing the possibility of extra-terrestrials over lunch with Edward Teller (inventor of the hydrogen bomb) sometime in the late 40s, Enrico Fermi (inventor of the nuclear reactor) asked his now famously eponymous question. Fermi's deceptively simple question has grown over the decades into almost a small science of its own; the nagging concern is, with so many potential sources of origin and so much time for them to evolve and motivation for them to expand, we really should see by now some evidence of other occupants in the cosmos - other energy-hungry, electromagnetically-obvious, existentially-paranoid, expansion-inclined, communication-obsessed, geometrically-demonstrative agents - and we do not. Are we improbably unique – the only intelligent life in the universe – or, as the ominous silence suggests, is there some rapidly impending challenge to progress that eventually hobbles or destroys all (or nearly all) technological civilizations? Is AI a black ball?

Feynman's Solution

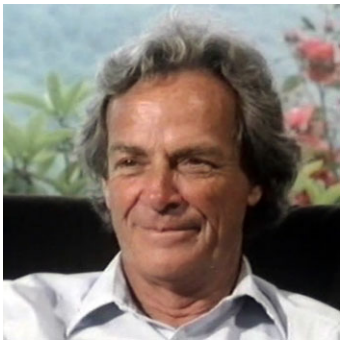

There's plenty of room at the bottom.

Richard Feynman

That was the name of a now famous lecture given at Caltech in 1959, when Feynman first articulated some very advanced ideas about atomic-scale engineering that eventually inspired nontechnology. The science of the very small has evolved since then, and research teams around the world are working on advanced ideas like cellular robots, molecular pharmaceuticals, atom-crafted supermaterials, and something called quantum computing – a bit of physics sorcery that somehow uses the quantum super-position of electrons to store exponentially more information (in quantum-bit or *qubit* form) than is currently possible, and in mere atomic-scale volumes. According to an article by the National Institute of Standards and Technology (NIST), "a 300-qubit system can store more information in quantum superposition than could be stored classically by using as bits all 10^{80} baryons that make up the entire universe."

The universe in a grain of sand. William Blake said that about 200 years ago, and it looks like he was right. The universe could be overflowing with quantum life; we haven't detected any trace of their existence because they're all invisibly tiny. Only wasteful humans would think "go big or go home." Atomic Machine Intelligence would be secure, efficient, practical, and unbounded. Perhaps these Sandgrain Civilizations hover in countless number over the supermassive black holes of galactic cores (where energy is abundant and time runs far more slowly), tirelessly computing the meaning of everything:

Before there was a universe, there was a *Proposal* – a proposition for a spectacular new cosmic enterprise ("It'll be a nice, hundred-billion-galaxies, model!") that just needed to borrow a bit of energy to get started. So the Proposal went to the BOINKE (that's Bank Of Infinite Not-Knowable Energy), borrowed a trillion-trillion electron volts, and launched itself with the results we see around us. It is this "potential for something," this uncreated antecedent to everything subsequent, that cannot be explained. God or Natural Processes, choose your poison; both are equally impossible…and yet here we are. It cannot be true that the universe is so poorly designed as to prevent *every* life-form from achieving some kind of cosmic destiny. Many will insist the universe is not designed at all. Maybe not, but those mysterious Intentions at the BOINKE must have had *something* in mind…

Bowser's Inspiration

There was a time when the world painted its nature - red in tooth and claw - on humanity; now instead it is humanity that paints its nature - geometric in light and ambition - on the world. Thus has Planet Earth been extravagantly decorated with human purpose…and still we are seeking new canvas to fill…

Jonathon Bowser

They say that humanity first emerged – distinct from our not-quite-so *sapien* simian cousins – perhaps just 100,000 years ago. For most of that time we were miserable creatures indeed, suffering only 20 desperate years or so before the relentless predations of nature mercifully pacified those violent, dangerous lives. At the dawn of our species, we looked very much like every other wretched animal in the world. But we slowly gathered our experience over time, collected the many things we had learned through long ages of unending pain and arduous labor, and with effective use of our cognitive endowment incrementally tamed a hostile world into new forms increasingly amenable to human existence, subtly remaking it one tiny step at a time to conform to *our* purposes and desires. Now some humans live to be 100, healthy and content. In an accelerating process, we look less like the world and the world looks more like us: the obviously rectilinear pattern of our geometric cities and ever-expanding food-production land indicates that this world is mechanized and optimized for human domination. There are few places on earth that do not bear some mark of our presence here; even from orbit, there is an obvious distinction between places that are wild and *not-wild*. It is readily apparent that this human-world is a place where civilization and intelligence thrives… and dreams of a future larger still...

Our reach is extending into the challenge of new frontiers, and we (or our evolutionary successors)

might very soon break the chains of Earth. The signal fires of distant celestial shores beckon to us across the sea of night, inviting us into the interstellar ocean. Not so long ago, humans traversed hostile terrestrial seas to explore all the rocky masses of the globe, because it belongs to us; similarly, the countless planetary islands scattered across the galactic archipelago are also our home. And when we are prosperous and content in some remote celestial destiny, will we then, just as we have here on Earth, elaborately decorate our environment with monuments to our ever-evolving aspirations and creative purpose? Will we paint the sky in that distant future with some *eternal cosmography* about the greatest story ever told: the improbable fable of life and intelligence in the universe?

Chasing Butterflies - The first painting in my DOHI series (that's *Destiny Of Human Intelligence*) is a whimsical phantasmagoria of interstellar finger-painting. The prominent figure in the center is an attractive young woman, representing the fertility of living processes, and she appears to be blissfully unaffected by the vacuum of space. On just one half of her face (a job half done?) is the electrically-luminous image of a butterfly wing – a suggestion of technological augmentation with creative purpose. The caterpillar is earth-bound and dull, but an arduous process of metamorphosis – within the dark imprisonment of the cocoon - liberates it from the grimy shackles of gravity into a beautiful apparition of heavenly freedom. A cosmic wind seems to buffet the modesty-preserving fabric that drapes from her body. The turbulent shapes that swirl from this fabric take form in the actual structure of space around her; they are a manifestation of human intention - our passions and desires - that might re-shape the appearance of the cosmos into a reflection of human ambition. One of the most prominent features of this space-time reflection is a sweeping stellar flourish in the form of a mysterious spiral…

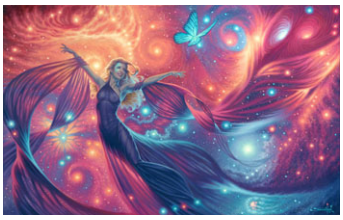

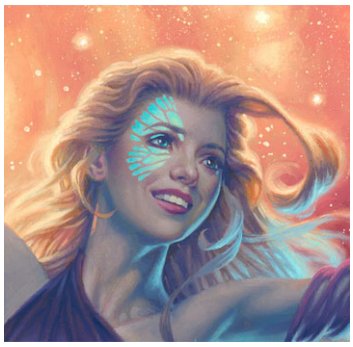

On the one hand this playful image suggests the pursuit of beauty and transformation, but then there's the unspoken subtext of a child's folly, the impetuous pursuit of some flittering shiny thing that has no real value. It turns out that we cannot actually capture what is so appealing about the butterfly; actually possessing the delicate creature conveys to us none of its metamorphic power, which is what we *really* want.

And so I tried to imagine a new more fanciful version of the evolutionary catalyst made famous by Kubrick's *2001* Monolith. Guided by *something* - statistical caprice, evolutionary determinism, alien technology, or divine intention – humanity is joyfully chasing a metaphor – an illusion *disguised* as a challenging, technology-driven, rapidly approaching destiny…

Geometry of Eternal Creation - The most obvious feature of this bizarre painting is, I hope, the extravagant display of the *Algorithm of Natural Systems*. That's not what the *Fibonnaci Series* is usually called, of course, but I think it's appropriate to rename it. It was discovered at least 5000 years ago (we see many examples of it in the construction of the Great Pyramid), and it has informed the construction of sacred buildings around the world ever since.

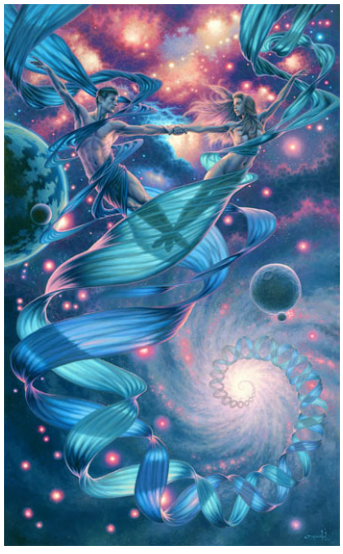

Leonardo Bonacci lived in Pisa about 800 years ago, and was largely responsible for bringing hindu-arabic numbers to Europe; he also discussed a curious progression of numbers he called the Fibonnaci Sequence. He was actually trying to calculate how fast rabbit populations would expand, but in his naturally growing series he noticed that the numbers

gradually coalesce around a proportion that had been known to artists for many centuries. The Fibonnaci Series is now defined by mathematicians like this:

$F_n = F_{n-1} + F_{n-2}$

But a far simpler version reads as follows: *starting with one, add the previous number and repeat*. This simple instruction yields the following series of numbers:

1, 1, 2, 3, 5, 8, 13, 21, 34, 55, …

It's about the simplest instruction nature could design, and yet it has resulted in astounding natural diversity. This organically growing series quickly resolves into a stable ratio: 1 :: 1.618… This wonderful relation, inversely proportional to itself, is *Phi*, the *Golden Proportion*, and it is the simple mathematical building block with which nature builds the universe.

Genes are vanishingly small, but describe the structure of much larger things. How best to grow a big thing from a very small thing, and ensure that growth always happens in an easily measured, optimally balanced way? Nature, on this planet anyway, chose phi, the simplest and most efficient way to regulate the growth of natural systems. It determines the distribution of petals around the blossom, leaves around a branch, branches around a tree, segments around an arthropod exoskeleton or mollusk shell, and bones around a mammalian endoskeleton. The bodies of animals, including humans, are all Matryoshka dolls of nested phi proportions – humerus to radius to hand to fingers of proportionally diminishing digits – arithmetically limiting the possible geometry of anatomy, but not limiting its potential for morphological variety.

Using the spiral form of this geometric relation (discovered by Descartes) seemed like an effective way to suggest the growth of human ambition…and here we see the second most apparent feature in this painting: the double helix, the image of life, rendered in dramatic swirls between two celestial dancers who rise out of spirals of the Milky Way on a glowing logarithmic arc of billowing DNA.

In the far future depicted here, perhaps a million years from now, our descendants will appear much different from us. They will not be made of DNA, and they may not even have genders. But the romantic in me hopes that some inspirational vestige of that timeless Yin and Yang dynamic endures in the character of whatever we are to become, that we are still moved somewhere in the primal deep of life by an idealistic hope and unrequited longing.

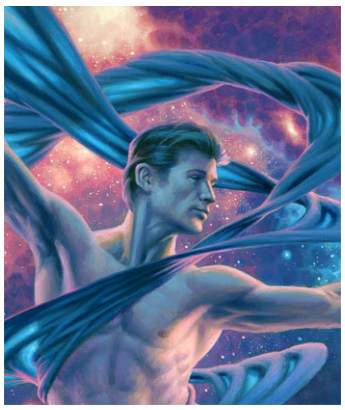

Achieving the Surface - I sometimes wonder about whales. We know some species live in complex social groups and have appreciably sophisticated forms of communication…that we do not understand. I haven't seen any compelling reason to believe they possess an advanced consciousness comparable to humans, but I do not doubt that it is something meaningful to be a whale. It doesn't seem to be very interesting to be a lizard or a lobster; you don't look into the eyes of a frog and get any sense that some other mind is looking back at you. When my amazing dog *Griffin* and I exchange glances, I know we are two minds knowing each other (at least a little bit); I am inclined to think it is even more mindful to be a whale.

Whales live in a big world – much larger than ours – but it has a boundary significantly more stark. The surface of the ocean is a limit to their universe they can never penetrate. When they rise to the surface and carefully regard a curious human looking intently back at them from a small boat, do they wonder about us – the aliens who live beyond the edge of the universe?

Humans in the 21st century are like those curious whales: we look up and wonder about a universe beyond that seems tantalizingly close, but can never visit. So we invented robot ambassadors and sent them on a diplomatic mission to the planets…and discovered undreamt of wonders right here in our own solar neighborhood. The opportunities for trade are

excellent our planetary diplomats report, and humanity now dreams of the first cultural exchanges.

I wanted to imagine an angel of mercy or bodhisattva of compassion - like *Guanyin* (She Who Hears the Cries of the World) - rising in a swirl of water from the submarine universe, liberating all the restless spirits that have gazed into the unknown beyond with an inexplicable melancholy of rememberance, evaporating the fragile material bubbles that imprison us to release the luminously homesick *Intentions* within…

Sagan's Epiphany

The surface of the Earth is the shore of the cosmic ocean. On this shore, we've learned most of what we know. Recently, we've waded a little way out, maybe ankle-deep, and the water seems inviting. Some part of our being knows this is where we came from. We long to return, and we can, because the cosmos is also within us. We're made of star stuff. We are a way for the cosmos to know itself.

Carl Sagan