

Rational Accelerationism

Archie Chaudhury June 2024

The day is November 29, 2022. The world of “tech”, comprised of snazzy journalists with RayBan glasses and comb overs, sweatshirt-draped college dropouts, product managers with inflated titles, Web3 connoisseurs with NFT profile pictures, and B2B SaaS founders working remotely on a Caribbean island, was collectively reeling from the potential financial windfall following the demise of public darling FTX while publicly decrying the Effective Altruism Movement and everything it stood for. 24 hours later, the so-called vibe had shifted, FTX was forgotten, and the aforementioned group of outcasts were collectively allocating their attention to an online chatbot, which, as fate would have it, will set forth perhaps the greatest technological innovation in human history.

Artificial intelligence has gone from being a far-fetched pipe dream that existed only in science fiction, to a legitimate field of study, to being used by millions of people all over the world in little over 60 years. OpenAI has been transformed from a non-profit research lab burning capital to one of the most valuable companies on the planet. ChatGPT has become a colloquial term, and is now a part of the vocabulary of many an over-caffeinated college student. AI has influenced movements that have led to the creation of new companies, reinvigorated entire industries, and caused monumental shifts in the fate of a few privileged companies and their fortunate shareholders.

Any discussion on AI will be without merit if it omitted the “doomerism vs accelerationism” debate. Various metrics, such as P-Doom, which rank different leaders by their predictions on the likelihood of a superintelligent AI destroying humanity, have been developed to chart the different positions on this debate. However, despite the innate philosophical attraction that comes with imagining a Detroit Become Human (play this game if you haven’t yet) future in which our metallic creations wrought of silicon rebel against us and lead to our ultimate downfall, it is not the biggest argument to be had in AI. Indeed, there is a larger, more influential battle brewing, one which very well may shape how this technology is developed, governed, and regulated.

Earlier this month, Leopold Ashenberner, a former OpenAI employee working on superalignment, published a set of long-form essays collectively titled “Situational

Awareness". In it, he argues that the potential development of Artificial Generalized Intelligence (AGI), which he predicts will be complete in less than 4 years, represents the largest national security hazard since the atomic bomb, and that in order to contain this threat, the US Government must ultimately nationalize the development of AI in order to protect crucial research from its enemies (China, North Korea, the usual suspects). Situational Awareness is a prediction of the not so far future, a future in which the development of AI has accelerated to the point that it is developing itself, a future in which the government must take control of AI or risk the equivalent of allowing "Sam Altman and Elon Musk to operate their own nuclear warheads". Its intended audience is not a member of government or nor the world of tech; rather, it is the users of AI, the stakeholders of a democratic government who will ultimately to some degrees decide how this technology is built and governed. Thus, each individual post is significantly more digestible than traditional technical works, and is meant to be read more of as an informal commentary (with some hard evidence) rather than a scientific essay only accessible by a few.

This work is meant to partially be a response/counter-argument to the ideology espoused in Situational Awareness, while simultaneously introducing a new line of philosophical thought that complements the existing "e/x" movements that live on the Internet: Rational Accelerationism. Despite the name, Rational Accelerationism is significantly different from the Rationalist movement popularized by Yudkowsky. Rational Accelerationism is rooted in the notion that the entirety of humanity, not just a small set of labs or government bureaucrats, can be trusted to move technology, science, and cognition forward in a responsible and ethical manner, that it should be independent, open-source scientists, not elderly statesmen, who should be on the forefront of governing artificial intelligence and other innovations. It is meant to be a summary of the outsiders, the ones who, while not being a direct part of the AI labs or the revolution in San Francisco, have meaningfully engaged with this technology for the better part of a year, and have been advocates for open-source, regulation-free development in all realms of technological innovation for even longer. This is an important discussion; while the accelerationists and their altruistic counterparts may very well be set in their beliefs, motives, and decisions, the broader public is not, and deserves to be, well-informed before a decision is made for them.

Table of Contents

This essay, like Situational Awareness, is organized into multiple independent subsections. These subsections can be categorized into three broader parts: part A delves into why it is rather unlikely for us to develop human-level intelligence relatively soon, part B discusses why it is unlikely for the US Government, or really any world power, to nationalize AI, and why unnecessary regulation may stifle any real progress we may have made on fundamentally improving the human condition, and part C introduces rational accelerationism, a movement advocating for the rapid, but safe, development of AI and adjacent technologies.

Each subsection can be read like an independent post. Part C has a short introduction, and then an independent manifesto introducing a new school of thought regarding the development of frontier technologies.

- A. We are still far from AGI*
 - 1. Numbers can lie: the reality of today's evaluation techniques and claims*
 - 2. The AI boom is not a boom, but rather a well-orchestrated explosion*


- B. Government Nationalization is unlikely*
 - 1. An unlikely parallel to the crypto wars of the 90s*
 - 2. Capitalism ultimately prevails for new technologies*

- B. introducing Rational Accelerationism*
 - 3. Safety is important, and should be in our hands*

A. We are still far from AGI

While the release of GPT-4 and other LLMs has taken the world by storm, it is important to recognize that fundamental limitations faced by the AI models we use today. The benchmarks and evaluations utilized to prove the cognitive capabilities of these models are flawed in their own right, and even if one is to believe them, suggest fundamental power laws suggest that we are probably further from generalized, and especially superintelligence, than what we may have been led to believe.

The release of GPT-3 to the public will without a doubt go down as one of the most monumental events in the interconnected histories of technology, science, business, and philosophical thought. Humanity was, without warning, presented with a strange reflection, a simulrca that presented itself like a human but was controlled by bits rather than flesh, powered by electricity rather than food. AI, once rooted mainly in the dreams of academics and 1960s science-fiction authors, was now in the hands of all. With opportunity came seekers, and soon enough, the technologist equivalent of the gold-rush began, with 100s of new models being released every week, both in an open-source and private manner.



Note the massive growth in foundational models produced by academia and industry, compared to the government, in this graph from a [report published by Stanford](#)

This new wave has prompted many an innovator, scientist, and thinker to proclaim that we are on the verge of AI Doom, a nightmare scenario in which humanity has to submit to its mechanical overlords due to our own irresponsibility and a lack of checks on the development of artificial intelligence. However, the reality today, outside from the miniaturized bubble of “tech”, suggests something else entirely, that AI, like the crypto bubble before it, is just that: a set of promising and innovative technologies that while having an extremely high potential, are nowhere close to fundamentally changing the human condition. AGI, or AI that reasonably perform at the cognitive level of an advanced human, is still far away, and more, not less, resources should be poured in making sure that we get there.

Numbers can lie: the reality of today’s evaluation techniques

The primary evidence that most “AGI is coming soon” proponents point to are the progressively higher scores achieved by frontier foundational models on popular benchmarks: in the [first chapter of Situational Awareness](#), Aschenberner cites the performance of GPT-4 on normalized educational aptitude tests, and FM-specific evaluation tests such as the MATH benchmark, a collection of medium to challenging mathematics questions with definitive numerical solution, as compared to its year-old predecessors.

MATH Dataset (Ours)

Problem: Tom has a red marble, a green marble, a blue marble, and three identical yellow marbles. How many different groups of two marbles can Tom choose?

Solution: There are two cases here: either Tom chooses two yellow marbles (1 result), or he chooses two marbles of different colors ($\binom{4}{2} = 6$ results). The total number of distinct pairs of marbles Tom can choose is $1 + 6 = \boxed{7}$.

Problem: The equation $x^2 + 2x = i$ has two complex solutions. Determine the product of their real parts.

Solution: Complete the square by adding 1 to each side. Then $(x + 1)^2 = 1 + i = e^{i\frac{\pi}{4}}\sqrt{2}$, so $x + 1 = \pm e^{i\frac{\pi}{8}}\sqrt[4]{2}$. The desired product is then $(-1 + \cos(\frac{\pi}{8})\sqrt[4]{2})(-1 - \cos(\frac{\pi}{8})\sqrt[4]{2}) = 1 - \cos^2(\frac{\pi}{8})\sqrt{2} = 1 - \frac{(1+\cos(\frac{\pi}{8}))}{2}\sqrt{2} = \frac{1 - \sqrt{2}}{2}$.

Professional Medicine

A 33-year-old man undergoes a radical thyroidectomy for thyroid cancer. During the operation, moderate hemorrhaging requires ligation of several vessels in the left side of the neck. Postoperatively, serum studies show a calcium concentration of 7.5 mg/dL, albumin concentration of 4 g/dL, and parathyroid hormone concentration of 200 pg/mL. Damage to which of the following vessels caused the findings in this patient?

(A) Branch of the costocervical trunk ✗


(B) Branch of the external carotid artery ✗

(C) Branch of the thyrocervical trunk ✓

(D) Tributary of the internal jugular vein ✗

Left: example of a question on the MATH Benchmark. Right: A question on MMLU, a comprehensive evaluation for LLMs.

However, there is a hidden problem with the majority of evaluations being utilized to prove the capabilities of most models used today: contamination. No, this is not a malignant virus or bacteria, but rather a phenomenon that describes how most evaluations end up inadvertently becoming a part of the very datasets used to train the models themselves, leading to those models overfitting on them and “memorizing” the answers to those evaluations. To visualize this, imagine that you are a high school student that has been gifted with an expanded hippocampus that can essentially allow you to memorize an unlimited amount of information at the cost of hampered reasoning ability. You can probably see how this allows you to essentially ace any test, as long as you have seen the questions and answers previously, but will result in you being stuck anytime you see a new question that requires you to apply the knowledge you have memorized. This is exactly what is happening with modern evaluations: the newest LLMs, which claim to be as prolific of a writer as Shakespeare while simultaneously being as mathematically inclined as Issac Newton, don’t hold up when presented with questions that were not public when they were trained.



Popular LLM performance on unaltered questions from MMLU (red) vs altered questions (everything else).
Graph from the following published [paper](#)

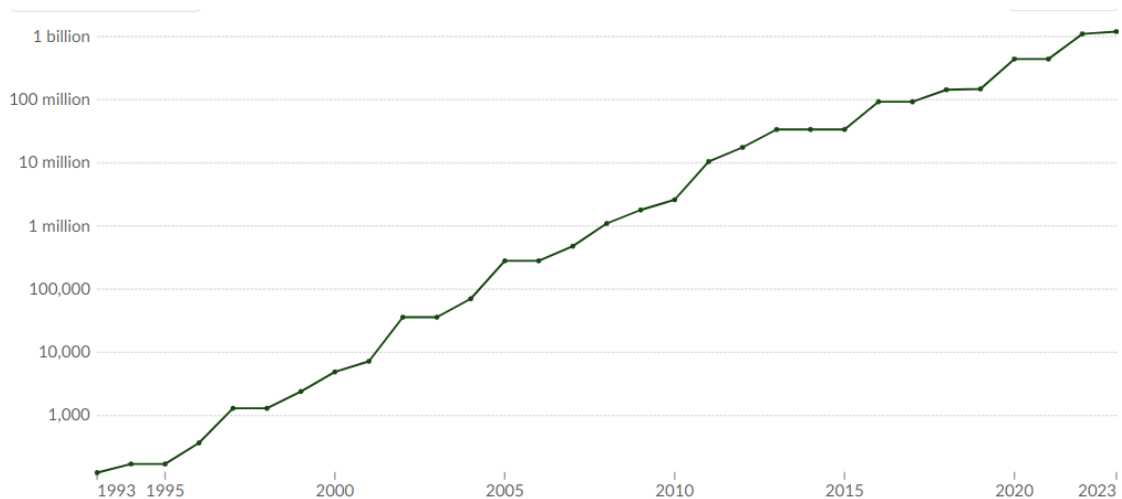
These results should not be astounding if you have interacted with LLMs in any meaningful way: while they are extremely useful at performing rote, computational tasks, they falter when confronted with something more complex. LLMs are extremely prone to hallucination, and often make critical mistakes. It is important to note that these phenomena are not inherently malicious; as Aschenberner himself notes, the proliferation of new models has far outpaced the creation of new benchmarks. Combine that with the fact that most large models are often training on a large subset of the

entire internet, and you can see why most of the public evaluations and tests out there are inadvertently being memorized by the newest foundational models.

Perhaps the best academic example of the progress, or lack thereof, that foundational models have made toward achieving anything even close to simulating cognition can be seen in the [Abstract Reasoning Corpus \(ARC\)](#), which is an evaluation for AI models designed to measure reasoning and complex skill acquisition developed by François Chollet, a researcher at Google. ARC works by challenging models to solve previously unknown tasks based on known examples, similar to an IQ test, rather than attempting to measure generalizable knowledge that may easily be acquirable. Most notably, contemporary models score far below human-level performance on ARC, indicating that while they may be able to perform rote tasks and operations at an extremely high level, they are still far behind the average human when it comes to reasoning and deriving solutions to new problems based on pattern recognition.

The AI “boom” is not a boom, but rather a well-orchestrated explosion


The central claim that most proponents of the “AGI soon” camp point rest their heads on is the progress made, both in terms of powerful hardware being readily made available to the public and highly efficient algorithms drastically reducing the learning speeds for foundational models, in the past couple of years. However, it is important to take a step back and deconstruct the phenomena that are TPUs and compute efficiencies. Moore’s Law is still very intact (and might even to some [extent be dead](#)) despite what you might have been led to believe by the accel mob on Twitter.



Graph from our world in data showcasing the growth in computational capacity of supercomputers, expressed in GigaFlops, over time

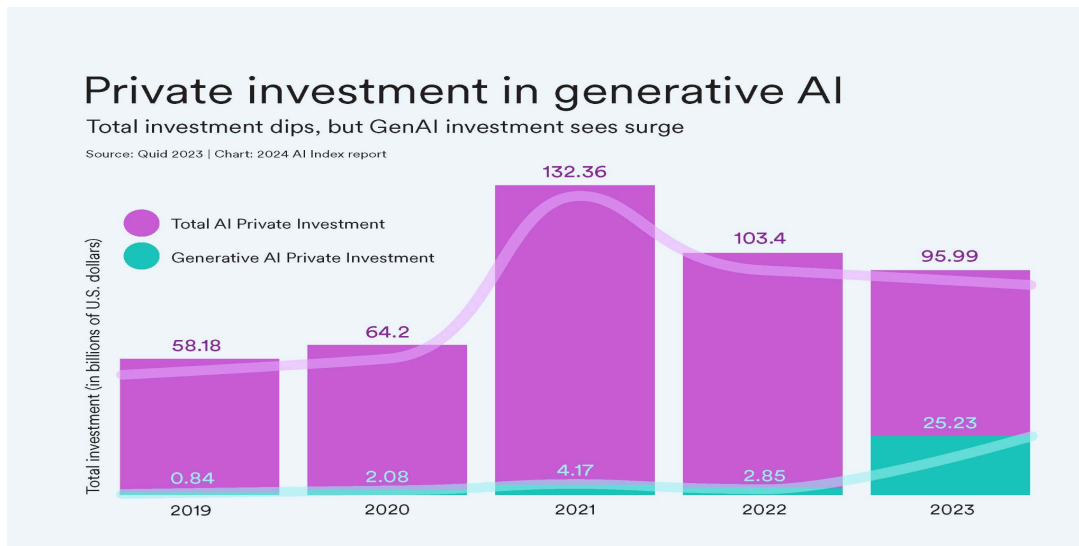
The recent growth that we saw for AI, after all, was not the product of some sudden, innovative hardware or algorithmic breakthrough, but rather the result of [years of](#) research, development, and experimentation. Indeed, capital inflow and the external attention given to generative AI has grown far beyond what one might have expected in recent years, but the capabilities of these models should not surprise anyone who was aware of the developments happening in the space pre-2022, especially someone who may have been a part of the San Francisco AI community that Aschenberner himself cites.

Now, you might be wondering: *Isn't that the entire point? That we need to accelerate the development of AGI/ASI?* And your claim will be right, but partially. The underlying crux of the arguments presented by Aschenberner and others who believe in an AGI-soon timeline (note that this includes both “doomers” and “accelerationists”) is that AGI is inevitable, and that we (the collective “we” meaning society, the US government, etc depending on who is making the argument) should be pricing the eventual development of AI into our lives now in order to get ready. However, at the current pace, *this is likely not true*. While improvements in algorithmic efficiency and hardware have certainly contributed to the development of GPT-4 and others of its ilk, the major contribution was funding: massive investments enabled frontier labs and established organizations such as Google to procure the hardware and talent needed to train monolith-sized models we see today. However, there is an upper limit to such improvements: for example, GPT-4, despite being significantly more aligned than its predecessor, was still as, and in some cases, even more, prone to hallucination.



Graph from Gorilla’s release of its LLM-tooling study

It is also important to note that the results of invested capital in frontier technologies often take years to manifest; indeed, the largest institutional investments in AI were during the 2020-2021 pandemic period, and are actually coming back down in recent years. Open-source AI has a [funding problem](#), with many academic labs and open-source projects falling behind larger companies. Assuming that the trend that we saw in 2022 and 2023, which was really the end of result of a multi-year R&D effort and a historic influx of capital, will not only continue but accelerate in an exponential manner is certainly a stretch.



Graph from Stanford Report showcasing investment in Gen-AI

It is important to note that the purpose of this section, and really the entire broader series that it is a part of, is not to present a gloomy-outlook on the current state of AI. Instead, it is to highlight the realities of how far we are from achieving AGI, and to implore the public collective to continue to move, tinker, and experiment rather than assuming that AGI is an inevitability. The current reality is that progress in fundamental compute has slowed down, not increased, and that algorithmic improvements, despite

improving the efficiency with which new models can be trained and developed, have not resulted in significantly increased cognitive performance for the publicly-available models that we use today.

B. Government Nationalization is unlikely


Assume that all the arguments made thus so far in this essay are false. Perhaps OpenAI has already achieved AGI internally, and is simply hiding it from the public, much like the government may be hiding sentient aliens. Or, there is indeed a breakthrough coming in hardware that will reverse the trend in hardware capacity, or new algorithmic research will improve compute to where everyone will have their own conscious AGI. Even if all of these things happen, it is highly unlikely that any government will pursue complete nationalization or centralization.

“A specter is haunting the modern world, the specter of crypto anarchy” - Timothy May in the Crypto Anarchist Manifesto

Much like the quote above predicts the rise of an anarchical movement resting on the development of cryptography, a fundamental tenet espoused in Situational Awareness is the idea that war is coming on the backs of AI, with the AI race soon becoming an arms race. Rather than the specter of cryptographic technologies, nation states are being faced with the reality that the most powerful technology created in the entirety of human history may end up outside their control, or worse, in the hands of their economic adversaries. Aschenberner makes an apt comparison to the nuclear arms race, suggesting that the potential advantage of whichever country reaches AGI first will have an advantage comparable to the advantage experienced by the United States post World War II.

However, this is under the assumption that a massive increase in productivity (which is ultimately the promise of software-based AGI) will be such a threat to human life that governments will be forced to not only accelerate, but regulate it in order to protect their own citizens and ensure that this power does not fall into the wrong hands. The truth is: *we have already lived through a similar phenomenon*. The proliferation of the internet and social media networks in the late 2000s and early 2010s was arguably

leveraged for political means far more than AI could ever be. Wars were no longer fought using steel and ammunition, but rather using misinformation campaigns and malicious hacks carried out by unidentified figures.



Cyber-war breakdown from Statista

Yet, the United States Government (USG) has not nationalized the social media, or the internet, at least not in the traditional sense. A similar phenomenon is likely to occur in AI.

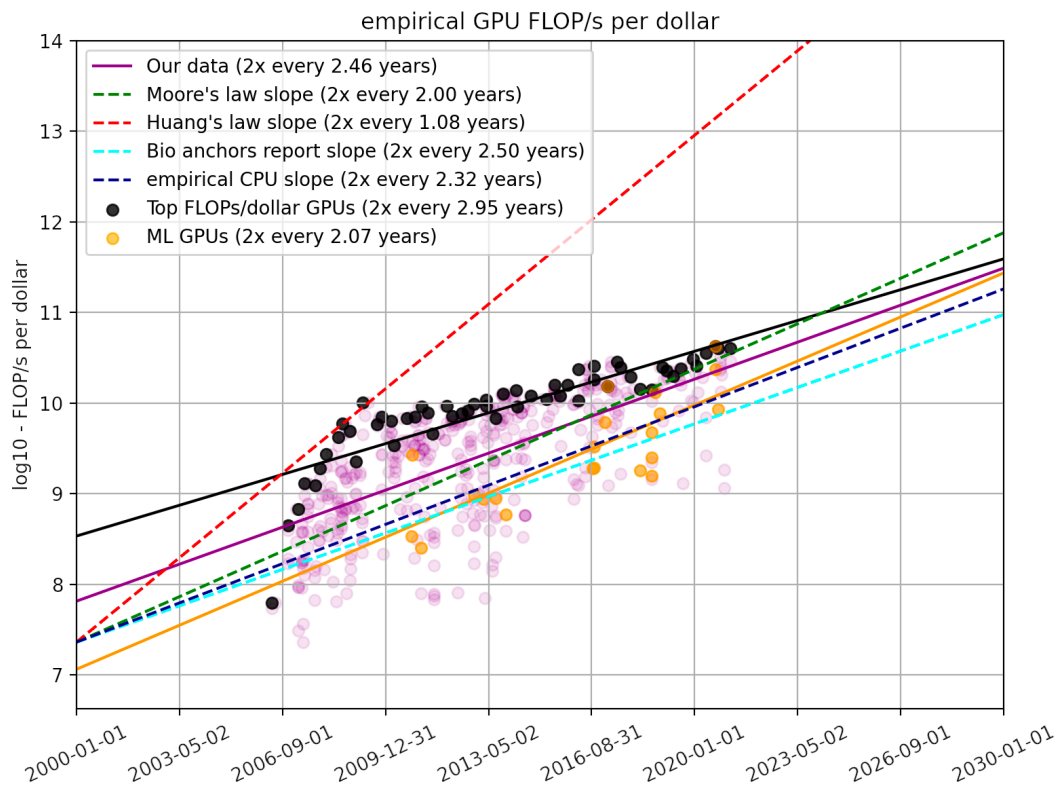
An unlikely parallel to the crypto wars of the 90s

The inclusion of Timothy May at the beginning of this particular section was not by accident; after all, May was one of the core figures in the cypherpunk movement that ultimately birthed Bitcoin, cryptocurrencies, and Web3. But, perhaps less famously, it was also the birthplace of many an [idea](#) in cryptography, privacy, and more. It is during this time period, in the early 90s, that the cypherpunk movement was the most prominent entity standing against the “war on crypto”, a proverbial term describing the USG’s attempts to restrict the export of privacy-preserving technologies to ensure that they did not fall into enemy hands.

Cryptographic protocols were soon declared munitions: creating and distributing a new privacy algorithm or anonymous email protocol online could carry the same charges as helping to move physical weapons over international borders. Perhaps notable victim of this was Phillip Zimmermann, the inventor of PGP, who over a three period was submitted to federal investigations by the USG. However, after a period of time, the USG eventually relented, both due to the difficulty of regulating what was essentially distributed open-source software and the impact that strict regulation had on for-profit entities (we will discuss why capitalistic needs almost always call for reduced regulation in the next subsection). Cryptography was, for the most part, free for use by all, and the government even stepped in to fund or develop open-source, privacy-preserving technologies. Tor, probably the most widely used privacy-preserving technology in the world, was developed by a group of independent computer scientists and researchers working for the USG Naval Research Laboratory.

So, what does this have to do with the possibility of frontier lab nationalization? Are these just the ramblings of a crazed free-market enthusiast? The key abstraction here is that foundational models are ultimately software, and software, like speech, is extremely difficult to regulate. The future in which a government such as ours is necessitated to nationalize or meaningfully impose its will on the development of AI is one in which both algorithmic and computational efficiency have increased to the point where the possibility of AGI is imminent. This is extrapolated from the main claims in Situational Awareness; AGI by 2027 is a possibility under the assumption that we get growth in hardware and algorithms that match, if not outpace, what we saw over the past three years. Yet, a future in which this occurs is also a future in which access to hardware is democratized, and most frontier algorithms that are published can be built and iterated upon by independent open-source scientists. Indeed, we can apply the same “straight-line” graph argument to show that if we assume hardware efficiency

continues increasing at the rate that it has maintained over the past 2 years (a rate of growth aptly termed “Huang’s Law”), then we will soon have highly-performant GPU clusters at a somewhat affordable price.



Graph from EpochAI predicting FLOP/s per dollar

In this scenario, hardware will be so ubiquitous that regulating AI will be seemingly impossible without causing substantial damage to larger businesses that will end up being the only ones to abide by them. It is hard to imagine a scenario in which the tech giants will be content with the open-source hobbyists ultimately outpacing them in AI, just like they weren't content with independent cryptographers outpacing them in implementing secure cryptographic protocols.

Capitalism ultimately prevails

More often than not, when it comes to new technologies, capitalism, and more specifically, the freedom to innovate, ultimately prevail. They are what allow one country to gain a technological lead over another, outweighing other aspects of economic production, such as the cost of labor or availability of natural resources. The United


States has consistently outpaced its rivals in technological development post World War II not because of its abundant talent or great natural industrial resources (though they have certainly helped), but because it has fostered a culture of innovation through incentives, comparatively regulation, and sparse opportunities for talented individuals with ideas or concepts to create companies of value. It is this form of incentive-based innovation that has allowed private companies and startups to outpace their academic contemporaries.

The USG has succeeded specifically because it has allowed new technologies to grow without imposing unnecessary regulations: AI is no different. This also applies to any strategic advantage that the U.S. may be seeking against its potential adversaries. Had the USG decided to regulate computers, the internet, social media companies, and search in the name of preserving national security, we may have seen China or some other authoritarian power become the default capital of the internet, and may be referring to Shenzhen rather than Santa Clara as Silicon Valley.


As Aschenberner describes in part three of *Situational Awareness*, a future in which the majority of AI research and development is done behind the closed doors of government-sanctioned facilities is one that will most likely lead to the US preserving its current academic/theoretical lead, at least for the short-term. However, this will soon dissipate. Modern frontier technologies are not developed in bubbles but rather through open, economic-based platforms. The majority of OpenAI's largest customers are not members of the public sector, but rather large companies and startups seeking to use AI in their own workflows. The release of ChatGPT in 2022 put forth an entire economy, mostly located in San Francisco, of random startups leveraging it for different use-cases. AI research has also become increasingly private; most of the world's foremost AI researchers are not working in academia, but rather on the research teams of large corporations such as Google, Meta, or Tesla. Indeed, the aforementioned three companies have arguably produced and done more for the development of technology over the past decade than the entirety of the USG put together (and the USG is happy about this!); it seems unlikely that curtailing their attempts to advance the state of AI research forward will lead to a positive outcome. In a manner similar to the development of another frontier technology, cryptocurrency, a large proportion of open-source AI developers are not constrained by national origins or boundaries. In fact, recent trends in the development of cryptocurrency and Web3 seem to pose the perfect example of what can happen when the USG decides to overregulate a particular industry: developers, businesses, and capital flows elsewhere where it is more likely to be accepted, and the U.S. as a whole falls behind.

The U.S. is losing its lead in web3

% of crypto developers in the United States



Top crypto websites* - % of traffic from the United States



Source: Electric Capital, SimilarWeb | *Includes Coinbase.com, Binance.com, Coinmarketcap.com, Coingecko.com, Etherscan.io.

Graph from A16z's state of crypto report in 2023

If AGI is to truly become an existential threat, then it is of utmost importance that we (again, we being the collective United States and its citizens) win the figurative race with our competitors to get there first. This is something I completely agree with *Situational Awareness* on: a world in which an authoritarian regime wields control over AGI is not a world in which the U.S. will reign supreme for long. However, the path there is not through nationalization and regulation; rather, it is through providing support and incentives to the innovators, the builders, the startups that are building and accelerating this technology. Aschenbrenner argues that the equivalent of allowing independent labs and startups to pursue AGI is the equivalent of allowing Uber to build nuclear reactors. However, we are no longer living in WWII: the public sector has fallen behind private development for decades, and now, even its strongest branches are reliant on technologies developed by for-profit companies (take the relationship between SpaceX and NASA for example). This model has enabled the U.S. to dominate the development of other technologies for the vast majority of recent technological cycles, and it will remain the same for AI.

C. Introducing Rational Accelerationism

It is important to note that some form of artificial general intelligence is likely to come if we continue expending a massive amount of both physical resources and talent toward AI research and development. However, accelerating AI does not have to come with a potential dystopia. In this final section, we introduce rational accelerationism, a new philosophy on the development of AI that probably has become the opinion of the silent majority.

Effective Accelerationism. Doomerism. Decentralized Accelerationism. "X" Accelerationism. There are a plethora of opinions regarding how humanity should approach the development of new frontier technologies. They, like the traditional socioeconomic spectrum of old, are marked by the schism between capitalism and socialism: on one end of the spectrum, effective accelerationists believe in unrestricted techno-capitalism, and on the other, "decelers" believe that technology will soon cause the death of humanity as we know it.

However, just like politics, philosophy, or economics, the vast majority of participants in AI are not extremists. They don't believe in technological advancement without any checks, nor do they believe that we need a proverbial big brother in the form of government regulations to make sure we don't develop something straight out *Terminator*. They are passionate about developing AI for the betterment of all, to eliminate manual labor and other tasks so that humanity can focus on creative tasks and expression. They believe that AI will one day be able to solve maladies and produce mathematical research that rivals that of our brightest, but also recognize that we are far away from needing to prepare for such a reality.

Rational Accelerationism is perhaps the best summation of the philosophical motivation behind this piece. *Situational Awareness* raises an interesting and pertinent point about the future of AI; Rational Accelerationism is a philosophy for why humanity, be it in the form of corporate scientists or anonymous developers with anime profile pictures, can be trusted to undertake in this future without the need for government oversight.

The following manifesto summarizes this school of thought, and at the same time, acts as a parting thought for this piece.

The Rational Accelerationist Manifesto

AGI is coming. There is no denying that. A small proportion of technologists are already readying themselves for a post-AGI future, focus on developing creative skills and output rather than intellectual or technological ability.

The rapid proliferation and adoption of AI has also raised numerous ethical questions about whether we are in reality creating a dystopia rather than a utopia that better the human condition as a whole. The most cautious of us have called for government-mandated pauses or oversight of AI development, taking it out of the hands of independent companies or startups, believing that civilizational collapse is in order if we are not responsible.

Yet, if there is one thing that has been made abundantly clear over the past century, it is that humanity can ultimately be trusted to produce abundance, to produce positive outcomes, when dealing with technologies. The net positive of the internet, which could have easily become a vector for unrestrained cyberwarfare and espionage, has far outweighed its negatives. Investment in nuclear energy and other forms of alternative energy generation have laid the groundwork for sustainable energy consumption. The development of space-exploration, which could have resulted in the loss of millions of dollars, has seen private corporations putting forth the vision for a future in which humanity becomes a multi-planetary species.

All of these advancements have come through individual corporations and technology firms operating independently, only relying on the government for support and guidance. The development of AI has the potential to not only be the technological movement with the largest potential to improve the human condition, but also one that can assemble the largest amount of independent talent. Engineers, scientists, policy experts and economists, for the first time, are all coming together to work on the same ideas.

From the anonymous developer you work with on Discord, to the entrepreneur doing the rounds on Forbes, we have shown our capacity to act rationally, to behave in a way that preserves us. Are we going to make mistakes as we head toward the development of AGI? Probably. Is that better than the alternative? **Without a doubt.** Safety and alignment are extremely **important to get right**: they are not just divisions within a company. However, the frontier labs currently leading the AI movement have not only shown a commitment to addressing such issues, but have spent actual capital on it. It should be up to the free-market, not our regulatory overlords, to guide to ensure that we build technologies that are both beneficial but safe. Our desire to survive is rooted in the most powerful instinct in the known universe, the human survival instinct, and will guide how we handle AGI, just as it guided the development of all the technologies that have gotten us here in the first place.