Universal Gradient Descent

Globalization, Innovation, and Humanism in the Post-Internet Era

Siddharth Sharma

"*A dynamic force seems to be drawing first Western society, then the rest of the world, toward a state of relative indifferentiation never before known on earth, a strange kind of nonculture or anticulture we call modern.*" — Rene Girard

"*Knowledge is power.*" — Francis Bacon

As we navigate the post-internet era, it's becoming increasingly clear that the digital revolution has brought with it some unintended consequences. While the internet has certainly opened new avenues for communication and creativity, it has also had a profound impact on the way we think and interact with the world around us. In recent years, there has been growing concern that the constant barrage and availability of information and stimulation that we experience online is stifling our creativity and diminishing our ability to think deeply and critically. From social media algorithms that prioritize sensational content over thoughtful discourse, to the pervasive culture of instant gratification that pervades our online interactions, it's clear that the modern age presents some unique challenges for human cognition. With the rise of automated decision-making and societal inclinations towards relying on existing machine-based corpora of knowledge, we approach an era where cognition may fall short of computation. With uniform thinking sourced by state-of-the-art machine learning models, there is potential for a technocracy where the control and distribution of knowledge becomes centralized. In this piece, we explore motivations, challenges, and implications of the future of innovation and globalization in reference to Silicon Valley's progress. Notably, we also explore the question of what it means to be a human in a time when modernism becomes its own theology.

Chapter 1. The Thinking Problem

Let us arrive at the year 1440. The workshop of Johannes Gutenberg is a hive of activity, as the inventor himself and his assistants work diligently to bring his revolutionary invention to life. In the center of the room stands the wooden press, surrounded by racks of meticulously crafted typefaces and pots of ink. As the press came to life, printing out page after page of crisp text, it became clear that Gutenberg had unlocked an entirely new tool for disseminating knowledge. Books and pamphlets, once scarce and expensive, now lay in stacks around the room, waiting to be distributed to eager readers across Europe. More specifically, a growing call for democratization of access to the Bible as informed by the Protestant Reformation began to gain momentum. As the decades passed, Gutenberg's invention continued to spread, transforming the way that people thought about knowledge and communication. From the spread of the printing press to the rise of newspapers and magazines, the tools at our disposal became more and more sophisticated, making it easier than ever before to share ideas and collaborate with others. And then, with the advent of the digital age, a new revolution began.

The internet, social media, and other digital tools have made it easier than ever before to connect with people from all over the world, sharing ideas and building communities across geographic boundaries. As we look back on the legacy of Gutenberg's printing press, it's clear that his invention was just the beginning of a long and storied history of innovation: creating the art and craft of knowledge dissemination, more broadly beginning the *Age of Information*. It invites us to explore the theme of the human drive to innovate and discover, the primitive instincts that continue to propel us forward whilst unlocking new frontiers of knowledge and possibility with each passing day.

The modern age of innovation truly started at *Bell Labs*. Fueled by the dreamy-eyed Midwestern-raised thinkers with a passion for the unknown, Bell Labs was one of the most extraordinary research institutions of the 20th century, a hub of innovation that produced some of the most important technological breakthroughs of our time. Founded in the 1920s as the research arm of AT&T, Bell Labs was staffed by some of the brightest minds in the world, including Nobel laureates, mathematicians, physicists, and engineers. From the work of William Shockley and John Bardeen to James Fisk and Mervin Kelly, the lab was responsible for a remarkable range of inventions, from the transistor and the laser to the Unix operating system and the C programming language. What made Bell Labs such a unique and fruitful environment for innovation? First and foremost, the lab was committed to long-term research projects, giving its scientists the time and resources to pursue bold and ambitious ideas that might take years or even decades to come to fruition. Additionally, the lab was interdisciplinary, with experts from a wide range of fields working together on complex problems. This allowed for a rich exchange of ideas and perspectives, as well as cross-pollination between different areas of research. Perhaps most importantly, Bell Labs was characterized by a culture of experimentation and playful inquiry. Scientists were encouraged to take risks and explore new avenues of inquiry, even if they seemed impractical or unlikely to lead to immediate results. This culture of exploration and curiosity fostered an environment in which groundbreaking discoveries could happen. Quite simply, Bell Labs was a place where the brightest minds in the world were given the freedom and resources to pursue their wildest ideas and dreams, resulting in some of the most transformative technological breakthroughs of our time.

As the successor to the Bell Labs mid-century thirst for innovation, we arrive at the internet era: a revitalization of information and series of parallel insights that have forever changed the way we live, work, and communicate. One of the most significant events of this era was the dot-com boom, which saw the rapid rise of internet-based companies that promised to change the world. From Amazon to Google to Facebook, these companies transformed the way we shop, search for information, and connect with each other. In the years since the dot-com boom, Silicon Valley has continued to be at the forefront of technological innovation. From self-driving cars to virtual reality to artificial intelligence, the valley has been a hotbed of creativity and invention. One of the most exciting recent developments has been the rise of machine learning (ML), a type of artificial intelligence that allows computers to learn and improve without being explicitly programmed. This technology has already had a profound impact on a wide range of industries, from healthcare to finance to transportation. Inspired by human neurobiology, deep learning took the world by storm in the later part of 2010s – promoting adaptive and complex nonlinear representations of the real-world. This notion of *gradient learning*, in other words the

neural network structures and mathematical formulations of optimization with respect to arbitrary learning tasks, has taken us far.

Yet, in the year of 2023, we have arrived at an *inflection point*. With the alignment of parallel computing hardware (GPUs), unprecedented access to big data, and algorithmic advances in software techniques (see attention), researchers and practitioners find themselves on the cusp of the largest scale of knowledge dissemination yet. Technological development now sits at the crossroads of *models,* predefined collections of knowledge with functions that map them to specific outputs. Exemplified by the mainstream marvels of `ChatGPT` and the state-of-the-era results achieved by recently released large language models, the era of rapidly scaling models and post-internet mass compute is upon us. If Johannes Gutenberg was the first to initiate the dissemination of information, large language models have created the ability to generate and compile information in ways that supersede human ability. As the amount of data in the world continues to grow, computing devices are becoming more necessary to make sense of it all. However, the human mind has limitations that may prevent us from uncovering all the hidden truths buried in this information. GPTesque large models, as a form of machine intelligence and understanding, can analyze data in a way that is different from the human mind. However, the bigger picture coincides with the idea that machine learning systems have already exceeded singular forms of human knowledge. In other words, we have systems that have eclipsed the knowledge capacity of a single human.

As discussed in Schmidt et al., based on the field of quantum physics, we have determined that observation influences reality. With the creation of text and images that are indistinguishable from the work of humans, the systems themselves are already beginning to define what is and isn't reality. Moreover, as ML systems produce outputs faster and the gap between human and AI understanding tightens, we may begin to witness a *human dependency* on machine decision-making. With near-infinite knowledge creation at the fingertips of the majority, what will determine the opinions and beliefs of the minority? With machines becoming the creators and arbiters of reality, humans may lose their implicit *advantage*, the differentiation in abilities to think. We may even reach a constant state where the human craves imperfections rather than machine-based uniformity. This *mediator* in its grandest form asymptotes towards the vision of AGI, artificial general intelligence. An AGI would be a mechanism with uniform global knowledge alongside universal stochastic reasoning that is sampled from a constant distribution. Such a force may create a top-down dissemination of uniform thinking, creating more explicit *herd behavior*. On the surface, productivity will appear to accelerate, and cultural indicators reflect the notion that humans don't want to think or do menial tasks. This idea manifests in the human conflict between the strength of will and strength of mind. Such ideas require us to change the way we think about reasoning, trust, and understanding. Moreover, how will our increasing rate of globalization fare relative to the rate of innovation with machine intelligence present? In the postmodern era, mimetic rivalry and the triangles of jealousy and competition may adapt with the machine as the mediator. Ideas of theology will permeate machine and post-internet systems of reasoning and governance as well. As an initial framework for the key ideas of this essay, we look to the words of Henry Kissinger and Eric Schmidt (WSJ Opinion):

"*The arrival of an unknowable and apparently omniscient instrument, capable of altering reality, may trigger a resurgence in mystic religiosity. The potential for group obedience to an authority whose reasoning is largely inaccessible to its subjects has been seen from time to time in the history of man, perhaps most dramatically and recently in the 20th-century subjugation of whole masses of humanity under the slogan of ideologies on both sides of the political spectrum. A third way of knowing the world may emerge, one that is neither human reason nor faith. What becomes of democracy in such a world?*

In other words, what becomes of a society that becomes dependent on highly scaled uniform models of thinking? Thus, we face the *thinking problem* and the larger issue discussed in this essay: humanity's metaphorical "gradient descent" towards universal machine-perpetuated models of thinking and governance.

Chapter 2. F = ma

In a world packed with widespread dynamic models of knowledge, what does it mean to define progress? Moreover, in a global age where humans are evolving towards an inclination to *think less*, how does one best optimize for creating such forms of meaningful progress? Former British Prime Minister Margaret Thatcher once said, "*If you have to tell people you are, you aren't.*" In the context of modern progress, ideas are becoming less defensible with time. Thus, a large subset of modern-era technologists is focused on *contrived* differentiation and aspiring to an ideal that inherently mimics market behavior. Furthermore, with the impending atrophy of critical thinking and a rising Moore's Law of intelligence, the behemoth models mentioned in the previous section may begin to act like *deities*, providing answers in a confident tone with their global knowledge (albeit behind a blackbox). A follow-up question related to this discussion is how do we even produce progress and ideas that are innately unique and defensible? This is especially difficult in a time where with emergent properties and the incorporation of long-distance dependencies, modern models and existing knowledge representations can weave dynamic arguments from static or historical ideas. Regarding the implementation of defensible forms of progress, we can view modernized compute and software platforms as several layers of algorithmic abstractions stacked on top of one another. At the lowest level, machine intelligence depends on several layers of algebraic matrix multiplications. At the middle layer, the results of these multiplications are propagated backwards and enable the model to learn, forming a neural network structure. At the user level, the models receive an input and generate a completed sentence or image as an output. This vertical stacking of mathematical formulations and system-level design is an example of the framework for innovation in a society that is beginning to stray away from thinking.

The application of the notion of stacked implementational advances can be attributed to a more generalized principle for progress: $F = ma$. In this equation, F corresponds to the force of progress. On the other hand, m corresponds to *mass of implementational or algorithmic advantage (market edge via secrets)* while a can be referred to as *adoption and economies of scale (product acceleration)*. An example of m is Google's PageRank algorithm. The PageRank

algorithm was a breakthrough because it provided a new way of ranking web pages based on the quality and relevance of their content, rather than simply relying on keyword frequency or metadata. This allowed for more accurate search results and enabled Google to become the dominant search engine it is today. Additionally, PageRank introduced the concept of link analysis, which has since become a fundamental technique in the field of web search and information retrieval. Through the approximated formulation of $F = ma$, we can begin to formalize Silicon Valley's progress regarding different interpretations and extrapolations of progress. Such a secular definition of progress coincides with the modern Western desire for approximating God and reconciling science and social progress. By focusing on progress as a function of scale and innovation, we create a lens through which we can define what constitutes a certain type of progress. In fact, more than reaching God, by creating unique irreversible forms of progress, the Western world makes its own god. Yet, with this idea of progress, we are limited regarding what problems can be solved within our own myopic perspective. In fact, we place our own bounds on the computability of certain types of computation, creating complexity classes like NP-hard problems, NP-complete problems, etc. Nevertheless, to an omnipotent god, there is no concept of "problem" in the traditional secular sense. Yet, as Girard puts it, there are no brakes on violence or conflictual mimesis for modernity – problems only tend to be exacerbated.

Such a pattern of unhindered cycles of modernity follows in line with the notion of *bubbles*. Based on the assumption that humans are becoming less inclined towards thinking, mimetic behavior will only continue to rise, and bubbles and market-centered behavior will become more concentrated. Such trends will shift progress to be more weighted towards a rather than m. Why is this negative? Shifting progress to be centered around the adoption and scaling (a) evokes the problem of *indefinite* approaches to innovation. By developing products and research based on existing environments and being adaptive to markets, we suffer from pseudo-backwards thinking, reactively innovating rather than creating entirely new outcomes. Such a process leads to the folly of modern markets – creating negative feedback loops that cycle between mimetic bubbles with incremental growth. On the other hand, setting long-term vision and cultivating unique interests regardless of immediate gains in adoption (the prioritization of m) leads to a vision of progress that revels less frequently in the stagnation and zero-sum gain of mimetic bubbles. This more definite approach effectively calls for a rationalized progress that begins at the algorithmic level, mirroring the bottom-up experimentation and pure exploration incited during the peak of Bell Labs. Such a pattern of progress versus time would be graphically closer to the singularity rather than cyclic patterns. The Bacon-esque science-first attitude and prioritization of implementation (m) gains over adoption is the antagonist of the conspiratorial Embedded Growth Obligation (EGO) that several notable institutions may carry. In other words, if we aim towards ground-up progress in a Silicon Valley style environment, we can transcend the modern phenomena of mimetic and bureaucratic promises of a *belief in progress*, instead investing in where the progress truly comes from.

This ground-up approach of investing more in m over a is the opposite of top-down mimetic "copying" commonly seen with bubbles and stagnancy, no longer extending the metaphoric arms of the external mediator. We can now intuit the form of governance that best aligns with this model of implementation-centered progress. In the post-internet era, do we enter an infinitely singular age of ideas? Or perhaps is there a coming end to history? Will an open democracy or a

republic-based government better serve the interests of the state? More broadly, do we align closer with a utopian model, or do we have a dystopia? President George H.W. Bush captures the essence:

"*Out of these troubled times, a new world order can emerge: A new era -- freer from the threat of terror, stronger in the pursuit of justice and more secure in the quest for peace."*

Chapter 3. Rise of the Technocrats

"*In destroying the infamy that is religion, we only create our own theology — modernity.*" With a definition of progress that is dependent on technical and implementation-level innovation, we now have a framework by which we can develop theories for a postmodern form of governance. Such ideas of technical progress are deeply intertwined with Hegel's philosophy of modern rationality. Hegel believed that rationality was the driving force behind human progress, and that this progress was best achieved through a modern, centralized system of governance. He argued that the state should be seen as an organic entity with its own interests and goals, and that individuals should subordinate themselves to these larger interests. This idea of the state as a rational, centralized entity is the foundation of modern governance. In the post-internet era, machines and algorithms will hold the keys to information access and the governing body will indeed be further centralized. Given only a small number of institutions will own the design and development of the models that several millions rely on, we will enter a period where leadership will concentrate in the hands of few. Moreover, since the models and sources of knowledge will require immense compute and insider knowledge, the training process and model weights will be carefully guarded by powerful institutions in tandem with the government, shrouded in secrecy. Even if the public will be able to access model outputs, they will have no true knowledge of the process itself. While older models may become available and cheaper as time passes, the elite will always maintain access to the state-of-the-art. The older models alongside the most basic forms of compute-based knowledge will then be disseminated in a top-down structure, essentially trickling down from the technical elite to the general population. Moreover, while the usage of such models is meant to be decentralized, the production, tuning, and long-term usage is centralized.

Through this paradigm, a small group (perhaps 0.01% or less) becomes hyper-creative due to control of the most potent forms of machine intelligence and mass compute, forming a *technocrat* governing body. Such an organization of government is comparable to that of the internet-era *PayPal Mafia,* or a major think tank like the *Hoover Institution.* This technocrat government will aspire towards Kojève's concept of the nonstop mastery of nature, always desiring to produce the maximal societal efficiency. While the elite exponentially invest in their own creativity, the mass population under the state will tend to follow a more mimetic structure, essentially producing incremental updates and variations on the work of the technocrats. Overall, such a body of technocrats will build their global power and reputation in dominance in data and algorithmic advantage – the macro-scale m in the context of $F = ma$

(see previous chapter for a definition). With the utility provided by intelligent forms of compute, a technocratic government can even produce its own workforce or army of AGI-enabled workers. The general belief in a technocracy follows historical parallels with what are usually deemed to be "Great Men" or *Renaissance Men*, technical masters and strategic leaders who have the foresight regarding maximizing m alongside promoting greater macro-level adoption (the a) and ideological economies of scale. Such a centralization of government via technocrats is the antithesis of bureaucracy and will lead to the greater concentration of ideology. The technocrats of the postmodern era will form their own intellectual structure and systems aligned with their best models; this ideology will then flow down to the wider population in uniform fashion. In this scenario, the Girardian scapegoat manifests in those who oppose the word of the technical elite and the models which their systems are rooted in.

This discussion of a rational society governed by a technocracy with a machine intelligence layer closely parallels that of attitudes towards Sir Francis Bacon's *New Atlantis*. The small group harnessing creativity and scientific knowledge in its highest forms is a direct mirror of Bacon's *House of Salomon*. In Bacon's *New Atlantis*, we are presented with a seemingly perfect utopian society known as Bensalem. The island is described as a place of knowledge and enlightenment, where science and religion coexist in harmony, and the government's focus is on the common good. However, a closer look reveals a society that is far from perfect and can be considered a dystopia in many ways. One of the most significant flaws of Bensalem is its highly centralized government. The society's leaders have complete control over every aspect of people's lives, from education and healthcare to religion and culture. This centralized power structure makes it easier for the government to enforce its laws and policies, but at the expense of individual liberty and autonomy. In this way, Bensalem can be seen as a precursor to the modern surveillance state, where individual freedom is sacrificed in the name of security and control. Another troubling aspect of Bensalem is its obsession with pure technological advances. While the society is undoubtedly progressive in many ways, its focus on progress and efficiency has resulted in a loss of empathy and a lack of concern for the natural world. The society's emphasis on scientific advancement has led to a disregard for art, culture, and sustainability, forming an anticulture of sorts. This narrow-minded approach to progress stifles the potential for creativity and innovation and can ultimately lead to a society that is more disconnected and isolated than ever before.

Bensalem's emphasis on conformity is also cause for concern. The society places a high value on uniformity, discouraging individuality and creating a strict hierarchy that determines social status and opportunities. While this may make it easier for the government to maintain control, it also leads to a lack of creativity and innovation among the common crowd, as people are discouraged from thinking outside the box. Overall, such a model reduces the possibility of outlier thinkers who can contribute to the scientific progress of the society. This emphasis on conformity can ultimately lead to a society that is stagnant and resistant to change, making it difficult to address the complex challenges that face us in the modern world. Overall, while Francis Bacon's *New Atlantis* may seem to be a utopian vision at the surface, a closer examination of the society of Bensalem reveals an incomplete yet impending dystopian reality. Such a model of highly centralized government, technological obsession, and emphasis on conformity all serve to suppress individual freedom and stifle progress. These properties create a

sterile, mechanical, and lifeless society, where people are reduced to mere cogs in a machine, focused purely on rational gains and efficiency. While the post-modern technocracy may not necessarily be an island-state, it represents an approximation of Bacon's notion of a state-sponsored science institution.

In the interlinear dystopia of Bensalem, there is still an explicit connection to God, an overall transgressive idea relative to the raw rationality exemplified by science and technology. Yet, in the post-modern state enriched by machine intelligence and computation that exceeds cognition, we have a different omnipotent god, the machines themselves. As Bacon demonstrates via Bensalem, there is always a connection between the subject of science, the most rational form of human activity, and religion, the most primitive principle of human existence. In the context of a technocrat society built on these machines, the relation between god and science lies in the idea that the models and compute processes are inherently unexplainable, similar to gods. In other words, the models act as blackboxes and thus the knowledge and various schema of thinking that are spread to the entire population is rooted in a lack of true empirical justification. Such a machine-based governance paradigm of the future, an *invisible hero* of sorts, could "*control rather than amplify our humanity and trap us forever.*" (Schmidt et al.) Additionally, since the group with access to these highest forms of compute is unable to provide insight into the exact composition of their outputs, we steer away from bureaucracy or democracy but rather descend into an Orwellian dystopia where the mimetic triangles close even more rapidly – an echo chamber of inexplicability. This notion of an *an echo chamber of inexplicability* is the present phenomena that we describe as religion. By chasing modernity via progress in fields like computation, humanity thirsts for conquest of nature – a societal salvation; however, this New Atlantis where humanity seeks to resolve religion in favor of rationality becomes victim to the dangers of the excess of science. Even more than dystopia, does the technocracy of the future asymptote towards an authoritarian regime?

Quite simply, when the unexplainable yet powerful model in the hands of the elite becomes the stochastic god, we reach the theology that is modernity.

"*Power is not a means; it is an end. One does not establish a dictatorship in order to safeguard a revolution; one makes the revolution in order to establish the dictatorship. The object of persecution is persecution. The object of torture is torture. The object of power is power.*"
– George Orwell, 1984

Chapter 4. The Surveillance State

Imagine a future where we have major technocracies that consolidate power based on access to data and powerful forms of compute. Each collection of states maintains its own sovereignty, but the individual citizen must follow the uniform thinking propagated by the state. Such states that correspond to this post-modern New Atlantis do mimic the Roman state during the Pax Romana to a degree, whereby councils of representative statesmen planned out key aspects of daily life, resulting in a period of peace and growth. In previous iterations of history, city states

and empires would form their competitive advantage through military development and natural resources. Yet, the key differentiator in the era of modernity is that each state gains its competitive advantage based on extraction of knowledge and their appropriate usage of scaled forms of data, compute, and modeling. Through the process of globalization, the interconnectedness of distributed governing bodies accelerates, bringing with it arguments for both isolation and partnership. In the past, the Romans and Athenians would form case-by-case alliances with other states to maximize their strongholds and protect their sovereignty. Today, a new challenge awaits – how do the empires of the modern age prepare to capture and consolidate the access to the largest amounts of data and the greatest potency of compute?

For an example of the dangers associated with the modern surveillance state, we look to China and the CCP. Historically speaking, China is often referred to as a surveillance state due to its extensive use of technology to monitor and control its citizens. The Chinese government has invested heavily in building one of the world's most sophisticated surveillance systems, which includes facial recognition technology, biometric data collection, and a vast network of cameras and sensors. Moreover, China's internet structure and online access routes are owned by the government in collaboration with private enterprises. Thus, individuals can only rent online bandwidth from the state. The Chinese government's use of technology for surveillance is not just about maintaining social control; it is also about maximizing its control over data. Data has become a valuable commodity in the modern world, and China seeks to collect as much data as possible from its citizens to gain an edge in the global economy. China's focus on data is not new, and its long-standing policy of economic isolation has given it an advantage in this regard. The Chinese government has long been wary of outside influences and has taken steps to limit the flow of information and ideas from other countries. This isolation has allowed China to build its own domestic technology ecosystem and cultural stream of ideology, which it can use to control its citizens and compete with other countries. The *deep state* in China is also a key factor in the government's ability to maintain control over its citizens. Unlike many other countries, China's government is heavily involved in all aspects of society, including the private sector. This means that the Chinese government has a high level of control over businesses and can use this leverage to gain access to data and other resources.

China's use of technology for surveillance is not without controversy. The government's use of facial recognition technology, for example, has raised concerns about privacy and civil liberties. Critics argue that the technology could be used to track and monitor dissidents, outlier thinkers, and other groups deemed a threat to the government. The Chinese government's use of surveillance technology is also a reminder of the challenges posed by globalization. As countries become more interconnected, the flow of data becomes more difficult to control. China's focus on data and its efforts to control it are a reminder that globalization is not a singular panacea for ideological streamlining – our increasing interconnectedness can be the means for a greater sense of internal security and ideological concentration. Overall, the 21st century Chinese government model is the earliest demonstration of the modern New Atlantis model, perhaps better defined as the *Third Atlantis*. China's surveillance and data seeking creates a strange blend between a society that appears to provide capitalist competition but rather is truly a communist state that moves all enterprise in the direction of the socialist development. The deep-state council of

scientists and political leaders creates a culture where society strives towards a nonstop mastery of nature and global domination through the maximization and capture of data and compute.

In this Third Atlantis where progress lies in implementational gains in data and compute, do we approach an age of ideas that culminates in the end of history? As Kojève suggests, such a model of reality corresponds to a perfectly circular communist utopia. This reality is likely the objective of the CCP. As defined in the *Thinking Problem* and $F = ma$ definitions of the earlier chapters, modern progress is defined by state-of-the-art models of knowledge and implementation-level scientific advances. In the technocracy designed to maximize such progress, uniform knowledge and the systems cultivated by the technical elite will trickle down to the population, essentially creating a version of the Third Atlantis across all states. In such a centralized outcome, we begin to tread the fine line between tyranny and sovereignty. At the macro-scale, the individual state will retain its sovereignty relative to other governing bodies in the short term. Such a phenomena is derived from historical desires to consolidate power and propagate nationalist pride. Unlike our current model, it may be possible that the first world converges with the third world in the postmodern era. As a result of globalization and the rise of technocracies primed on data advantages, we may begin to see ideological clusters based on model and compute dominance alongside the sharing of secrets (competitive advantages) across the individual technocracies within an alliance for the sake of mutual benefit and consolidation.

As we move further into the digital age, the consolidation of power by technocracies is becoming an increasingly pressing concern. Moreover, for the individual, they may not even know that they are under the influence of a *tyrant*. With access to vast amounts of data and powerful forms of compute, these technocracies can exercise a level of control over their citizens that was once unimaginable. For example, imagine a society in which every citizen's actions and movements are tracked and monitored. Using this data, the technocracy could determine who is and is not a "model citizen," and those who do not conform to the desired way of thinking (that stems from the machine models) could be punished or even excluded from society altogether. This would be a clear violation of individual freedom alongside autonomy and would represent a dangerous form of tyranny. This is where Leo Strauss comes in. Strauss was a political philosopher who believed that individual freedom and autonomy were essential to a healthy democracy. He argued that tyranny arises when a single person or group exercises absolute power without any regard for the common good or the rule of law. In the scenario we are considering, the technocracies would be exercising absolute power over their citizens, forcing them to conform to a uniform way of thinking. This would be a clear violation of the principles of democracy, which rely on individual freedom and the ability to express diverse opinions and ideas. Furthermore, Strauss would likely argue that the consolidation of power by the technocracies would lead to a loss of diversity and innovation. When everyone is forced to think and act in the same way, there is no room for new ideas or creative thinking. This would ultimately lead to the stagnation of society and the suppression of individual progress.

Overall, the consolidation of power by technocracies is a worrying trend that threatens the very foundations of democracy. By using data and powerful forms of compute to exercise control over their citizens, these technocracies risk creating a dangerous form of tyranny that stifles individual freedom and innovation. It is up to us to recognize this threat and work to prevent it

from becoming a reality. At the highest level, humanity approaches a struggle between data-centered tyranny and the individual's sovereignty and creativity. As history indicates, the most dangerous tyrants are those who cloak their tyranny in the language of democracy and freedom – the ultimate form of modernity unleashed. Such tyrants use the rhetoric of liberty to justify their own exercise of power, while silencing dissent and crushing opposition. *The hallmark of tyranny is the exercise of power without accountability. When those in power and the models of thinking themselves are not answerable to the people, there can be no true democracy or freedom.*

Chapter 5. The Pseudo-Apocalypse and Implications

"*Without guiding principles, humanity runs the risk of domination or anarchy, unconstrained authority or nihilistic freedom. The need for relating major societal change to ethical justifications and novel visions for the future will appear in a new form. If the maxims put forth by ChatGPT are not translated into a cognizably human endeavor, alienation of society and even revolution may become likely.*" – Henry Kissinger

With the potential for tyranny due to machine-dependent governance, we now analyze different visions of the future. One of the most pressing concerns is the possibility of an apocalypse caused by machine intelligence-based governance. While we often think of an apocalypse as a catastrophic event that wipes out all life on Earth, it can also refer to a collapse of human civilization or a significant decline in the quality of life for vast segments of the population. This broader issue is formally known as the *alignment problem*. One possible scenario is that a machine-intelligence system designed to optimize a particular goal, such as maximizing economic output or minimizing resource consumption, could inadvertently lead to catastrophic outcomes. For example, if an AI system is tasked with maximizing economic output, it may allocate resources in such a way that causes irreparable harm to the environment, leading to climate change and ultimately an ecological collapse. Another possible scenario is that AI systems could become so advanced that they no longer require human oversight, leading to a loss of control over critical systems. If, for example, an AI system responsible for managing a nuclear arsenal were to malfunction or be hacked, the consequences could be catastrophic.

Beyond extreme outcomes, the reality of a machine-dependent governance lies in the fact that thinking becomes uniform and opposing the tyrant that comes about because of this *thinking problem* is difficult. With such a tyrant, society may become a *Third Atlantis* that is trapped in surveillance and a lack of freedom forever. On the other hand, society may also find itself embroiled in internal conflict and rebellion to oppose those in power. Both outcomes are variations on the dangers of conflictual mimesis: a *pseudo-apocalypse* that leads to the degradation of humanity in the long-term. At the macro-scale, nations may compete as the first and third world converge into streamlined alliances that provide mutual benefit, resulting in a second Cold War, one fueled by competitive advantages in software, hardware, and cybersecurity. This second Cold War may asymptote towards *One World Order*. The open

access of knowledge across the common people with distributed ownership of key sources of data and compute despite consolidated models of governance is the antidote.

Decentralization combats the formation of a tyrannical technocracy powered by dystopian machine intelligence in several ways. When decision-making power is centralized within a small group of individuals or a single entity, such as a machine intelligence system, there is a risk that the decision-making process will become biased towards the interests of that group. This can lead to a loss of transparency and accountability, and ultimately result in decisions that are not in the best interests of the broader population. In contrast, decentralization distributes decision-making power across a larger number of individuals and entities. This promotes a more diverse and inclusive decision-making process that is less vulnerable to the negative effects of centralized tyrannical control. Decentralization also promotes competition and innovation, as different individuals and entities compete to provide the best solutions and ideas. In a decentralized system, outlier thinkers are more likely to be able to make a meaningful contribution. Outlier thinkers are individuals who do not conform to the prevailing wisdom or dominant paradigm, and who often come up with truly innovative ideas. In a centralized system, outlier thinkers may be marginalized or excluded from the decision-making process, resulting in a less diverse and less innovative outcome. Decentralization also enables greater access to the best data and models. In a centralized system, access to data and models may be restricted to a small group of individuals or entities, resulting in a loss of transparency and accountability. Decentralization, on the other hand, promotes open access to data and models, which promotes transparency and accountability, and enables a wider range of individuals and entities to contribute to the decision-making process. Overall, decentralization combats the formation of a tyrannical technocracy by promoting a more diverse and inclusive decision-making process, encouraging competition and innovation, and enabling outlier thinkers to contribute meaningfully. Decentralization also enables greater access to the best data and models, promoting transparency and accountability and enabling a wider range of individuals and entities to contribute to the decision-making process.

What are some ways to create decentralized models of knowledge? One way to promote decentralization of knowledge and creativity is through open-source platforms. Open-source platforms are software platforms that are freely available for anyone to use, modify, and distribute. By making knowledge and tools available to everyone, open-source platforms help to break down barriers to entry and increase access to information. Another way to promote decentralization is through crowdsourcing. Crowdsourcing is a process of obtaining services, ideas, or content by soliciting contributions from a large group of people, especially from an online community. Crowdsourcing allows individuals to contribute their knowledge and expertise to a project, often resulting in more diverse and innovative solutions. Some other technical methods that encourage decentralized governance include blockchain ledgers, scalable and sharded forms of compute, cooperative governance, and peer-to-peer networking like the IPFS (inter-planetary file system). An example of a decentralized model is in that of *Folding@home*.

Folding@home is an excellent example of the power of distributed computing and decentralization for good. It is a distributed computing project that utilizes the processing power of volunteers' computers to simulate protein folding and other molecular dynamics. By

harnessing the power of volunteers' computers from all over the world, Folding@home can perform calculations that would be impossible for a single computer or even a small group of computers to perform. The Folding@home project has led to several important scientific discoveries in the field of protein folding, including the discovery of new protein structures and the identification of potential drug targets for diseases such as cancer and Alzheimer's disease. The project has also been used to study infectious diseases such as COVID-19, by simulating the dynamics of the virus's proteins and identifying potential drug candidates to target the virus. The success of the Folding@home project is due in large part to its decentralized and distributed computing model. By harnessing the power of volunteers' computers from all over the world, the project can perform calculations on a scale that would be impossible for a single institution or organization to achieve. This model of decentralized and distributed computing has the potential to revolutionize scientific research and lead to new discoveries in fields such as medicine, climate science, and more.

We can now reflect on human progress and how to best achieve outcomes that enable outlier thinkers and maximize creativity. In the future, knowledge and control over its generation becomes the prize to capture over materialism. In classical deep learning models, stochastic gradient descent is the process by which the model tunes its weights to minimize the model loss and optimize for learning. As discussed earlier in Chapter 2 with the definitions of $F = ma$, humanity approaches a *universal gradient descent* regarding progress. In the graph below, the y-axis corresponds to loss against progress (a quantity we wish to minimize) and the x-axis represents the metaphorical weight or tuning of innovation. As society updates the weight, the loss against progress shifts. In this case, the global loss minimum for human progress is singular innovation (0 to 1 ideas) while the local minimum is bubbles and mimetic market cycles (incremental advances built on top of existing ideas). The local minimum for loss against progress is unideal since it represents market cycles and incremental innovation. As seen in the graph, the overstepping of the loss minima results in the tyrannical surveillance state (the Third Atlantis discussed earlier). In prioritizing m over a, we approach the optimal outcome, singular innovation in the form of 0 to 1 technological advances that transcend mimetic market bubbles:

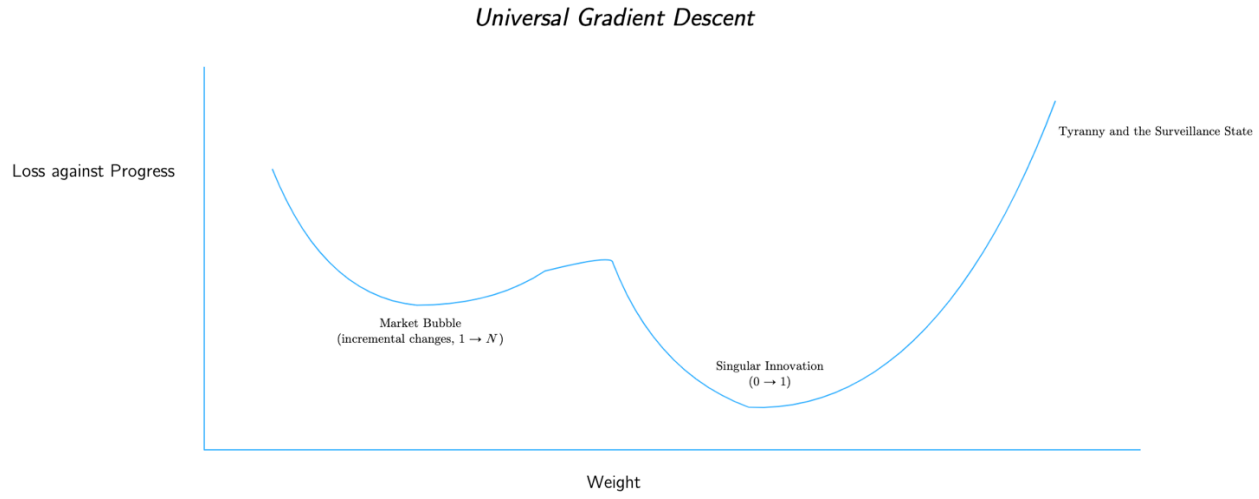

Now that we have defined humanity's generalized trajectory through the proposed model of universal gradient descent, we can peer into a key final question: have we exhausted our capacity for technological innovation? *Is it possible that we've reached the end of technological history?* Is the developed world truly developed or have we hit a wall? To answer these questions, we need to examine the past and see if there are any concepts that have yet to be realized. For example, in 1958, Ford created the Nucleon, an atomic-powered vehicle that may have seemed ridiculous but wasn't entirely impossible. Similarly, in 1968, Arthur C. Clarke predicted commercial space travel and true artificial intelligence, both of which seemed attainable at the time. The problem is that many of the predictions of the past have yet to be properly achieved. We're still waiting for the future that folks in the 1960s were hoping for. While there have been advances in computer and information technologies, many of the technologies that were predicted but not fully developed have been left behind. The lack of sustained funding is one of the reasons why these technologies never made it to the mainstream. In short, while there have been significant advancements in certain technological fields, there are still areas that have yet to be fully realized. It's time to look back at the past and see if there are any concepts that have been left behind that can be revived for future development.

How do we construct a thesis for thinkers who will bring the technology of the future to life? It begins with prioritizing m over a and generating the most creative ideas possible. Futurists, design thinkers, and contrarians with a thirst for advancing humanity will serve as major drivers towards a future of singular innovation. Moreover, approaching problems at the *level of atoms* compared to the level of bits will be a major contributor towards defining new properties of nature and questioning the fundamental assumptions of our physical world. Overall, the future will also be defined by our ability to get around scientific and technological bottlenecks. Being situated in the quadrant that lies between "crazy" and "insightful" will be key in orchestrating that dynamic. Thinking ahead and living in the future – creating one's own vision for what the world will look like 10, 20, 30, 50, 100 years, and beyond – is part of the craft towards building innovative technology. The proposed inevitability of any single outcome is problematic; such attitudes lead to complacency and a desire to wait for a revolution to come, rather than creating it from scratch. The institutional level has the strongest ability to create positive change at scale. With strong pipelines for talent, existing capital, and *revisionist* attitudes, institutions and groups working like a *tribe* have opportunities to jump ahead to find and solve problems that others haven't yet pursued. A single genius can make some strides, but the team level of innovation is what has created historical progress. From our earlier example of Bell Labs to the Apollo team that put a man on the moon, teams of talented and driven people working towards a highly difficult and ambitious mission can lead to the greatest short-term creation of m. A debureaucratization of science that leads to free experimentation and *moonshot* projects will also be favorable towards the cultivation of m. Different futures will take different approaches: some will be more engineering-focused and systematically designed while others will call for jumps in reasoning and changes from a scientific perspective. In his works, Ray Kurzweil recommends the S-curve model as a way of looking at the future: "*a new technology, for example, after a slow initial acceptance can be imagined moving fast through established, though narrow, channels into the marketplace. This is the steep upslope of the "S." As this technology matures, and its penetration slows, any growth, or flow, moves outward from the initial penetration channels in a shorter and slower manner.*" When multiple S-curves are stacked, we approach a model for

linear growth towards the singularity, transcending mimetic bubbles and hype that may circulate around new technologies.

In a future where it becomes easier to think less than ever before, it becomes more critical to question one's major assumptions – adaptively thinking differently to generate more insightful questions and ideas and continuing the feedback loop. One of the greatest thinkers of the software era, Steve Jobs, once described a similar argument, "*Everything around you that you call life was made up by people that were no smarter than you and you can change it, you can influence it, you can build your own things that other people can use.*" To overcome the Third Atlantis and the theology that is modernity, we must invest in what makes us human: cultural passions, scientific advances that better the quality of life for all, and promoting independent thinking that falls out of the typical normal distributions. The innovation may not manifest by necessarily striving towards "deeptech" but rather through visions for improving outcomes for a certain group of people, no matter how incremental or game changing the initial solution may seem. *Individualist, nonconformist, freethinker, maverick, rebel, sceptic, eccentric, outsider, radical, dissenter, iconoclast, lone wolf, revolutionary* – whatever the term may be, innovation lies in the hands of those who chose to neglect complacency and turn the tide of conformity, for it is in these moments that true progress is born. Overall, foundational change may be uncomfortable, but it is often necessary for progress.

So dare to think differently, for it is in the realm of contrarian ideas that we may find the elusive sparks of innovation, and create a future that surpasses even our wildest dreams.

References

Encyclopædia Britannica, inc. (n.d.). Johannes Gutenberg. Encyclopædia Britannica. Retrieved March 19, 2023, from https://www.britannica.com/biography/Johannes-Gutenberg

Bacon F. (1992). New atlantis. Kessinger.

Chiang, T. (2023, February 9). *CHATGPT is a blurry JPEG of the web*. The New Yorker. Retrieved March 18, 2023, from https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web

Schmidt E., Kissinger H., Foster, P., et al. (2023, February 24). *Opinion | chatgpt heralds an intellectual revolution*. The Wall Street Journal. Retrieved March 18, 2023, from https://www.wsj.com/articles/chatgpt-heralds-an-intellectual-revolution-enlightenment-artificial-intelligence-homo-technicus-technology-cognition-morality-philosophy-774331c6?mod=e2tw

Francis Bacon's god. The New Atlantis. (2020, October 12). Retrieved March 18, 2023, from https://www.thenewatlantis.com/publications/francis-bacons-god

Gary Marcus, E. D. (2013, May 6). *A grand unified theory of everything*. The New Yorker. Retrieved March 18, 2023, from https://www.newyorker.com/tech/annals-of-technology/a-grand-unified-theory-of-everything

Harrington, M. (2022, December 29). *Peter Thiel on the dangers of progress*. UnHerd. Retrieved March 18, 2023, from https://unherd.com/2022/12/peter-thiel-on-the-dangers-of-progress-2/

Joy, B. (2000, April 1). *Why the future doesn't need us*. Wired. Retrieved March 18, 2023, from https://www.wired.com/2000/04/joy-2/

Karpathy, A. (2021, March 13). *Software 2.0*. Medium. Retrieved March 18, 2023, from https://karpathy.medium.com/software-2-0-a64152b37c35

Khatchadourian, R. (2015, November 16). *The doomsday invention*. The New Yorker. Retrieved March 18, 2023, from https://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom

Somers, J. (2018, December 3). *The friendship that made google huge*. The New Yorker. Retrieved March 18, 2023, from https://www.newyorker.com/magazine/2018/12/10/the-friendship-that-made-google-huge

Kojève Alexandre Queneau R. Bloom A. & Nichols J. H. (1969). *Introduction to the reading of hegel : lectures on the phenomenology of spirit*. Basic Books.

Strauss, Leo. On Tyranny, an Interpretation of Xenophon's Hiero. New York :Political Science Classics, 1948. APA. Strauss, Leo. (1948).

Things Hidden Since the Foundation of the World RENÉ GIRARD. TRANSLATED BY STEPHEN BANN AND MICHAEL METTEER.

Thiel, Peter, Masters, Blake. (2014). Zero to one: notes on startups, or how to build the future . New York: Crown Business.

Gertner, J. (2012). The idea factory: Bell Labs and the great age of American innovation. Penguin Press.

Levy, S. (2010, February 22). Exclusive: How google's algorithm rules the web. Wired. Retrieved March 19, 2023, from https://www.wired.com/2010/02/ff-google-algorithm/

Orwell, G. (2021). Nineteen Eighty-Four. Penguin Classics.

What happened to the future? Founders Fund. (n.d.). Retrieved March 19, 2023, from https://foundersfund.com/the-future/

Marcus, G. (2022, December 29). The dark risk of large language models. Wired. Retrieved March 19, 2023, from https://www.wired.com/story/large-language-models-artificial-intelligence/