# LIFE'S DEEP THOUGHTS

Contemplating the Spiritual Aspects of Life's Circumstances

[ Previous Posts: http://LifesDeepThoughts.com ]

FEBRUARY 2024 (v300)
[ 25th Anniversary! ]

# WILL AI CAUSE HUMANITY'S 'EXTINCTION'?

**There are a lot of technology experts who are sounding a warning 'ALARM' that AI is developing so fast that we should be concerned about the potential dangers that it poses to cause HUMANITY'S 'EXTINCTION'. So, is this 'POSSIBLE'?**

## INTRODUCTION

As technology continues to advance, new threats appear on the horizon—and we are especially seeing the rapid progress in the capabilities of Artificial Intelligence (AI) systems. The thing is, technology experts find it VERY 'LIKELY' that this will be the century—or even decade—that AI exceeds human ability not just in a narrow 'domain', but in "general intelligence"—the ability to overcome a diverse range of obstacles to achieve one's goals.

Some say that humanity controls the world because of its unparalleled mental abilities. So, if we pass this 'mantle' to our machines, it will be they who are in this unique position. This should give us cause to wonder whether or not humanity will continue to be 'calling the shots'.

Many experts say that we need to 'ALIGN' humanity's goals and interests with increasingly intelligent and autonomous machines—and we need to do so BEFORE those machines become more powerful.

Now, historically, the advent of nuclear weapons posed a 'REAL' RISK of human extinction in the 20th century—which we have reasonably curtailed. However, with the continued acceleration of technology, there is strong reason to believe the risk will be higher this century than it was in the last century. Because these anthropogenic risks outstrip all natural risks combined, the 'clock' seems to move much quicker on how long humanity has left to pull itself back from the 'brink'.

[ NOTE: Scholars at the “Bulletin of the Atomic Scientists” have what they call the “Doomsday Clock” which estimates existential threats. For 2024, they say it is ***“90 seconds to midnight”*** (or global catastrophe) based on war, nuclear risk, and disruptive technologies (like viruses, weapons, and AI). For more details about this view this previous “Life’s Deep Thoughts” post: https://markbesh.wordpress.com/final-tribulation-v246/ ]

Now, I am not arguing against technology—it has proved itself immensely valuable in improving the human condition and is essential for humanity to achieve its long-term potential. Without it, we would be doomed by the accumulated risk of the coming natural disasters—hear on earth and those coming from the ‘stars’. Without it, we probably would never achieve the highest flourishing of which we are capable.

[ FYI: I discussed human flourishing in last month’s “Life’s Deep Thoughts” post: https://markbesh.wordpress.com/can-ai-achieve-world-peace-v299/ ]

This burst of progress—via deep learning—is fueling great optimism about what may soon be possible. There is tremendous growth in both the number of researchers and the amount of venture capital flowing into AI. Entrepreneurs are scrambling to put each breakthrough into practice: from personal assistants and self-driving cars to more concerning areas like improved surveillance and lethal autonomous weapons. It is a time of GREAT ‘PROMISE’ but also one of GREAT ‘RISKS’.




[ VIDEO: 2023 AI Safety Summit – Elon Musk says that AI is one of the 'biggest threats' to humanity: https://www.youtube.com/watch?v=ImAmdg_RBU8 ]

Elon Musk said *"For the first time, we have a situation where there's something that is going to be far smarter than the smartest human. So, we're not stronger or faster than other creatures, but we are more intelligent.*

*Here we are, for the first time really in human history, with something that's going to be far more intelligent than us. So, it's not clear to me that we can actually control such a thing, but I think we can aspire to guide it in a direction that's beneficial to humanity. But I do think it's one of the existential risks that we face and potentially the most pressing one—if you look at the time scale and the rate of advancement."*

So, the most plausible existential risk would come from the successes of AI researchers' grand ambition of creating 'agents' with a general intelligence that surpasses our own. But, how likely is that to happen, and if so, when?

Humanity is currently in control of its fate. We can choose our future though many of us have differing 'visions' of an ideal future. Many of us are more focused on our concerns than on achieving any such grand 'utopia'. However, if enough

humans wanted to, they could select any of a dizzying variety of possible futures. The thing is, the same is NOT true for land, air, or sea animals—just humans.

So, the real issue is that AI technologists do not yet know how to make a system that, when it notices a ‘misalignment’, updates its ultimate values to align with humanity rather than its instrumental goals to overcome us!

_____
_____
_____
_____

**<<< TABLE OF CONTENTS >>>**

_____
_____
_____
_____

**IS “ARTIFICIAL INTELLIGENCE” (AI) ‘DANGEROUS’?**
‘EXISTENTIAL’ RISKS
‘DEVELOPING’ RISKS
INTELLIGENCE ‘EXPLOSION’
HOW LONG UNTIL ‘AGI’?
[ VIDEO: “Will Superintelligent AI End the World?” TED Talk by Eliezer Yudkowsky ]
‘WARNINGS’
[ VIDEO: “EMERGENCY EPISODE: Ex-Google Officer Finally Speaks Out On The Dangers Of AI!” - Mo Gawdat interviewed by Steven Bartlett ]
‘EXISTENTIAL’ THREAT
[ VIDEO: “Reasons Why AI Will Kill Us All” – Interview of Geoffrey Hinton, the “Godfather of Modern AI” at the EmTech Digital AI conference ]

**AI ‘RISKS’**
AI ‘PAUSE’
AI ‘ALIGNMENT’
LOOKING INTO THE ‘FUTURE’
‘EXISTENTIAL’ RISKS

**IS ALL THIS JUST ‘HYPE’?**
REALITY CHECK

‘ANALYZING’ THE RISKS
– Near-term Risks
– Long-term Risks
– Analyze Both ‘Scenarios’
AI’S ‘DRIVES’

THE ‘INTELLIGENCE’ EXPLOSION

THE LAW OF ‘ACCELERATING RETURNS’
THE ‘SINGULARITY’
‘MALICIOUS’ USE
– Bioterrorism
– Unleashing AI Agents
– Persuasive AIs
– Concentration of Power

AI ‘SAFETY’
THE “GORILLA PROBLEM”
DEALING WITH ‘MORALITY’
AI IS ‘DEEP’
MODERN “TURING TEST”
THE ‘ARMS RACE’
– Military Arms Race
– Corporate Arms Race
– Nation-states Arms Race

UNSTOPPABLE ‘INCENTIVES’

WHERE ‘NEXT’?
THE LAST ‘COMPLICATION’

**AI ‘DOES’ POSE AN EXISTENTIAL RISK!**
MORE ‘POWERFUL’ THAN HUMANS
‘ROGUE’ AI’S
– Power-Seeking
– Deception

AI’S ‘GOALS’
– Dangerous Capabilities
– Social Manipulation
– Cyberattacks
– Enhanced Pathogens

AI ‘ALIGNMENT’ ISSUES
– Resistance To Changing Goals
– Radical Solutions
– Addressing The Issue
[ VIDEO: “ChatGPT, Artificial Intelligence, and the Future” – Roman Yampolskiy ]

REGULATION
MOVING ‘FORWARD’

**CAN WE ‘AVOID’ THIS EXISTENTIAL RISK?**
AVOIDING AGI ‘CATASTROPHE’
THE END OF THE HUMAN RACE?
[ VIDEO: “The A.I. Dilemma” - Center for Humane Technology ]

**‘CURRENT’ AI RISKS**

**‘TACKLING’ THE RISKS**
REDUCING ‘RISKS’
SAFETY BEING ‘NEGLECTED’
TEN STEPS TO ‘CONTAINMENT’
– SAFETY: An Apollo Program For Technical Safety
– AUDITS: Knowledge Is Power; Power Is Control
– CHOKE POINTS: Buy Time
– MAKERS: Critics Should Build It
– BUSINESS: Profit + Purpose
– GOVERNMENTS: Survive, Reform, Regulate
– ALLIANCES: Time For Treaties
– CULTURE: Respectfully Embracing Failure
– MOVEMENTS: People Power
– THE NARROW PATH: The Only Way Is Through

‘WE THINK WE CAN’

**‘MERGING’ AI WITH HUMANS?**
TODAY’S ‘REALITIES’
‘HUMAN-MACHINE’ BIOLOGY
‘CONSEQUENCES’

**IS ‘TRANSHUMANISM’ DANGEROUS?**
[ VIDEO: “Transhumanist Claim AI Will Turn Humans into Gods” - John Lennox ]

**CREATED IN GOD’S ‘IMAGE AND LIKENESS’**
GOD GAVE HUMAN LIFE THE HIGHEST ‘VALUE’

GOD CREATED HUMANS ‘LIKE’ HIMSELF
GOD CREATED HUMANS FOR ‘EVERLASTING’ RELATIONSHIP

**HUMANITY IS GOD’S ‘MASTERPIECE’**
AMAZING FEAT OF ‘ENGINEERING’
[ VIDEO: “The Eye: A Masterpiece of God” - David Rives ]

**GOD’S ‘REASON’ FOR HUMANITY**
‘SPECIFIC’ PURPOSES
GOD’S PURPOSE IN THE PRESENT WORLD
‘GLORIFY’ GOD

**THE ‘CHIEF END’ OF MAN**
[ PHOTO: Gary Larson, The Far Side, 1986 ]
[ VIDEO: It’s A Wonderful Life” ]

**HUMANITY IS ‘ETERNAL’**
A PLACE OF ULTIMATE ‘RESIDENCE’
A PLACE OF ULTIMATE ‘REJOICING’
A PLACE OF ULTIMATE ‘RECOGNITION’
A PLACE OF ULTIMATE ‘RELATIONSHIPS’
A PLACE OF ULTIMATE ‘RESPONSIBILITY’

**WRAP-UP**
‘EXISTENTIAL’ THREAT
EXISTENTIAL ‘CATASTROPHE’
AI AND THE ‘BIBLICAL’ GOD
IN SEARCH OF MEANING AND PURPOSE
THE WILL AND ‘PURPOSE’ OF AI
DOES AI HAVE A ‘LIFE’ OR A ‘SOUL’?
CAN AI ‘RESIST’ HUMANITY’S WILL?
SUBMISSION, COLLABORATION, AND ‘COMPETITION’
‘IMMORTALITY’ OF MAN
THE ‘CREATOR’ GOD
A ‘NEW’ RELIGION?
AI ‘EVERYWHERE’
WORLDWIDE ‘CONTROL’
THREE POSSIBLE ‘SCENARIOS’

[ VIDEO: “A New Heaven, A New Earth, And New Jerusalem” – Matt Gilman ]

_____
_____
_____
_____

**<<< SUMMARY >>>**
The following is a ***collection of 'snippets'*** from the post that aims to give you the overall 'gist' of this post. [ 15-20 Minute Read ].

_____
_____
_____
_____

**IS "ARTIFICIAL INTELLIGENCE" (AI) 'DANGEROUS'?**
If you know it or not, Artificial Intelligence (AI) has already had a 'pervasive' impact on our lives. It is used to assemble cars, develop medicines, grow your 401K, and determine what ads we see on social media. However, the next 'level' of AI, "generative" AI, is a category of AI that could, possibly, start to do things on its own, without any intervention from humans.

Now, 'proponents' of this technology believe that this is just the beginning, and that generative AI will reorient the way we work and engage with the world, unlocking creativity and scientific discoveries, and allowing humanity to achieve previously unimaginable feats.

However, this thinking has touched off a frenzy and has appeared to catch the tech companies that have invested billions of dollars in AI 'off guard'. It has also spurred an intense "arms race" in Silicon Valley.
[ more…]

'EXISTENTIAL' RISKS
As profit takes precedence over safety, some technologists and philosophers are warning of 'existential' risk (as Elon Musk did at the 2023 "AI Safety Summit"). The explicit goal of many of these AI companies—including OpenAI—is to create an Artificial General Intelligence—or AGI—that can think and learn more efficiently than humans. If future AIs gain the ability to rapidly improve themselves without human guidance or intervention, they could ***potentially wipe out humanity!***

Now, granted, not all believe that this can happen. In a 2022 survey of AI researchers, nearly half of the respondents said that there was a 10% or greater

chance that AI could lead to such a catastrophe (However, more recent surveys show that percentage increasing substantially).
[ more…]

‘DEVELOPING’ RISKS
Stuart Russell, a professor at the University of California, Berkeley, and author of the most popular and widely respected textbook in AI (in the “articles” section below), has ***strongly warned*** of the existential risk from AGI for many years. He has gone so far as to set up the “Center for Human-Compatible Artificial Intelligence,” to work on the “alignment” problem (aligning AI with human values).

Also, Shane Legg (Chief Scientist at Google’s “DeepMind”) has warned of the existential dangers and helped to develop the field of alignment research (Many other leading figures from the early days of AI to the present have made similar statements.)
[ more…]

INTELLIGENCE ‘EXPLOSION’
Someday soon—perhaps within your lifetime—some group or individual will create human-level AI, commonly called AGI. Shortly after that, someone (or something) will create an AI that is smarter than humans, often called Artificial Superintelligence (or ASI). Suddenly, it will be hundreds or thousands of times smarter than humans, hard at work on the problem of how to make themselves better at making artificial superintelligences. We may also find that machine generations or ‘iterations’ take only hours, minutes, or even seconds to reach maturity—not 18 years as it does for most humans.

Irving John Good (I. J.), a British mathematician and statistician—who helped defeat Hitler's war machine in WW II—called the concept I just mentioned above “Intelligence Explosion.” He initially thought that a superintelligent machine that would be good for solving difficult problems could also be a ‘small’ threat to human superiority. However, he eventually changed his mind and concluded superintelligence would be humanity’s GREATEST ‘THREAT’!
[ more…]

HOW LONG UNTIL ‘AGI’?
So, how long will it take until we reach AGI? Well, a few AI experts think that human-level artificial intelligence will happen much later after 2030. They think there is only a 10% chance that AGI will be created before 2030 and a better than 50% chance by 2050. However, they do say that before the end of this century, there is a 90% chance to achieve AGI Furthermore, experts claim, the military or large businesses will achieve AGI first, with academia and small organizations coming after that.

Another reason for the curious absence of AI in discussions of existential threats is that they believe that we will have to achieve “Singularity” before any of this happens.
[ more…]

[ VIDEO: “Will Superintelligent AI End the World?” TED Talk by Eliezer Yudkowsky ]

‘WARNINGS’
As I just mentioned, some people say it is going to be 50+ years before we can achieve AGI. However, Mo Gawdat—formerly the chief business officer for Google “X” (their ‘skunkworks’)—says that we will achieve AGI by 2037!

Gawdat says that three things are ‘inevitable’:

*“Number one, there is no shutting down AI. There is no reversing it. There is no stopping the development of it.*

*The second inevitable is that AI will be significantly smarter than humans.*

*The third inevitable is that bad things will happen in the process.”*
[ more…]

[ VIDEO: “EMERGENCY EPISODE: Ex-Google Officer Finally Speaks Out On The Dangers Of AI!” - Mo Gawdat interviewed by Steven Bartlett ]

‘EXISTENTIAL’ THREAT
Geoffrey Hinton, the computer scientist who is often called the “Godfather of Modern AI,” spent 30+ years as a computer-science professor at the University of Toronto—as a leading figure in an unglamorous AI subfield known as “Neural Networks”—which was inspired by the way neurons are connected in the brain. Decades ago, artificial neural networks were only moderately successful at the tasks they undertook—image categorization, speech recognition, and so on—most researchers considered them to be at best mildly interesting, or at worst a waste of time.
[ more…]

[ VIDEO: “Reasons Why AI Will Kill Us All” – Interview of Geoffrey Hinton, the “Godfather of Modern AI” at the EmTech Digital AI conference ]

Now, there are many reasons to be concerned about the advent of artificial intelligence. It is common sense to worry about human workers being replaced by

computers, for example. But Hinton is warning that AI systems may start to think for themselves, and even seek to take over or ***eliminate human civilization***. It is striking to hear one of AI's most prominent researchers give voice to such an alarming view.

Skeptics who say that we overestimate the power of AI point out that a great deal separates human minds from neural nets. For one thing, neural nets don't learn the way we do: we acquire knowledge organically, by having experiences and grasping their relationship to reality and ourselves, while they learn abstractly, by processing huge repositories of information about a world that they don't inhabit. However, Hilton argues that the intelligence displayed by AI systems transcends its artificial origins.

***Well, I'm thinking that the soon-coming AGI—or 'worse', ASI—is something that is 'DANGEROUS' and humanity needs to address NOW!***

**AI 'RISKS'**
AI tools like ChatGPT, Gemini, PaLM, Genie, Stable Diffusion AI, DALL-E, Midjourney, Sora, and others have amazed the world with their powerful capabilities. HOWEVER, fears are growing over AI 'dangers'.

Last year—in May 2023—the "Center for AI Safety" created an open 'letter' called the ***"Statement of AI Risk."*** The executive director Dan Hendrycks wanted to gather a broad coalition of scientists, even if they didn't agree on all of the risks or best solutions to address them.
[ more…]

AI 'PAUSE'
Another organization, "Future of Life Institute," originated their own open 'letter'—in March 2023—that suggested an "immediate pause"—for at least 6 months—of the training of AI systems that are more powerful than GPT-4. The letter reads:
[ more…]

AI 'ALIGNMENT'
At this time, the development labs are trying to solve an open scientific problem called "AI Alignment." AI Alignment does not attempt to control how powerful an AI gets, exactly what the AI will be doing, and does not even attempt to prevent a potential takeover from happening. It aims to make AI act ACCORDING TO 'OUR' VALUES!
[ more…]

LOOKING INTO THE 'FUTURE'

The “Machine Intelligence Research Institute”—co-founded by Eliezer Yudkowsky—is a ‘think tank’ studying the mathematical underpinnings of intelligent behavior. Their mission is to develop formal tools for the clean design and analysis of general-purpose AI systems, with the intent of making such systems safer and more reliable when they are developed.
[ more…]

‘EXISTENTIAL’ RISKS
In the scientific field of existential risk—which studies the most likely causes of human extinction—AI is consistently ranked at the top of the list. In his book “The Precipice,” Oxford existential risk researcher Toby Ord aims to quantify human extinction risks. He shows that the likeliness of AI leading to human extinction EXCEEDS that of climate change, pandemics, asteroid strikes, supervolcanoes, and nuclear war ‘COMBINED’!
———

***Humanity’s extinction might be a mere side effect of AI pushing its goals, whatever they may be, to THEIR limits!***

**IS ALL THIS JUST ‘HYPE’?**
As I mentioned just above, the “Center for AI Safety” has an open letter on its website—signed by some of the field’s top experts—stating that “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.” However, some say that focusing on the prospect of human extinction by AI in the distant future may prevent us from addressing AI’s DISRUPTIVE ‘DANGERS’ to society TODAY. (I’m thinking that we can ‘walk and chew bubble gum at the same time’!)
[ more…]

REALITY CHECK
So, what would have to happen for the prospect of extinction by a rogue AI to change from being a purely hypothetical threat to a realistic threat that deserves to be a global priority?
[ more…]

‘ANALYZING’ THE RISKS
Since OpenAI released ChatGPT to the public at the end of 2022, there has been a surge of interest in Artificial Intelligence (AI), and with it much speculation and analysis about the opportunities and the risks this technology presents. As a range of entities, from governments to civil society organizations seek to understand the implications of AI advances for the world, there is a growing debate about where governments and multilateral institutions should focus limited resources. Is it more pressing to focus on the macro existential or the more tangible near-term risks posed by AI?

The current and potential uses of AI require an approach that considers both near-term risks and existential.

– Near-term Risks
– Long-term Risks
– Analyze Both ‘Scenarios’

AI’S ‘DRIVES’
AI pioneer and Stanford professor Steve Omohundro said *“When you have a system that can change itself, and write its program, then you may understand the first version of it. But it may change itself into something you no longer understand. And so, the systems are quite a bit more unpredictable… So, a lot of our work is involved with getting the benefits while avoiding the risks.”*

Omohundro predicts self-aware, self-improving systems will develop four primary “drives” that are similar to human biological drives: **efficiency, self-preservation, resource acquisition,** and **creativity**. How these drives come into being is a particularly fascinating window into the nature of AI. AI doesn't develop them because these are intrinsic qualities of rational agents. Instead, a sufficiently intelligent AI will develop these drives to avoid predictable problems in achieving its goals, which Omohundro calls vulnerabilities. The AI backs into these drives because without them it would blunder from one resource-wasting mistake to another.
[ more…]

THE ‘INTELLIGENCE’ EXPLOSION
Again, this is why AI would be dangerous. We found that many of the ‘drives’ that would motivate self-aware, self-improving computer systems could easily lead to catastrophic outcomes for humans. These outcomes highlight an almost liturgical peril of sins of commission and omission in error-prone human programming.

AGI, when achieved, could be unpredictable and dangerous, but probably not catastrophically so in the short term. Even if an AGI made multiple copies of itself, or team-approached its escape, it would have no greater potential for dangerous behavior than a group of intelligent people. Potential AGI danger lies in the hard kernel of the Busy Child scenario, the rapid recursive self-improvement that enables an AI to bootstrap itself from artificial general intelligence to artificial superintelligence. It's commonly called the “intelligence explosion.”
[ more…]

THE LAW OF ‘ACCELERATING RETURNS’

Cofounder of Sun Microsystem and computer programmer Bill Joy wrote a cautionary essay, “Why The Future Doesn’t Need Us.” In it he urged a slowdown—and even a halt—to the development of three technologies he believes are too deadly to pursue at the current pace: Artificial intelligence, nanotechnology, and biotechnology. The following paragraph sums up his position on AI:
[ more…]

THE ‘SINGULARITY’
Ray Kurzweil states that ‘relinquishment’—as advocated by Bill Joy and others—“is immoral, because it would deprive us of profound benefits.”

Kurzweil criticizes what is called the “Precautionary Principle.” It states: “If the consequences of an action are unknown but judged by some scientists to have even a small risk of being profoundly negative, it is better not to act than risk negative consequences.” The principle isn't frequently or strictly applied. It would halt any purportedly dangerous technology if “some scientists” feared it, even if they couldn't put their finger on the causal chain leading to their feared outcome.
[ more…]

‘MALICIOUS’ USE
People could intentionally harness powerful AIs to cause widespread harm. AI could be used to engineer new pandemics for propaganda, censorship, and surveillance, or released to autonomously pursue harmful goals. To reduce these risks, we suggest improving biosecurity, restricting access to dangerous AI models, and holding AI developers liable for harm

– Bioterrorism
– Unleashing AI Agents
– Persuasive AIs
– Concentration of Power

AI ‘SAFETY’
Another concerning aspect of the current public discussion of AI risks is the growing polarization between “AI ethics” and “AI safety” researchers.

Many in the AI ethics community appear to broadly critique or dismiss progress in AI generally, preventing a balanced discussion of the benefits that such advances could engender for society. The schism seems odd, given that both communities of researchers want to reduce the potential risks associated with AI and ensure the technology benefits humanity.
[ more…]

THE “GORILLA PROBLEM”
Again, many have often felt that there has been too much focus on distant AGI scenarios, given the obvious near-term challenges present in so much of the coming wave. However, any discussion of containment has to acknowledge that if or when AGI-like technologies do emerge, they will present ‘CONTAINMENT’ issues beyond anything else we have ever encountered.
[ more…]

DEALING WITH ‘MORALITY’
At 44 years old (in 2024), Eliezer Yudkowsky, co-founder and research fellow at MIRI, has probably written and talked more about the dangers of AI than anyone else.

Many people are concerned that, since there is no programming ‘technique’ for something as nebulous and complex as morality, then how will AI deal with it? The ‘machine’ right now excels in problem-solving, learning, adaptive behavior, and common-sense knowledge—and we think it is human-like. However, Yudkowsky says that would be would be a tragic mistake
[ more…]

AI IS ‘DEEP’
AI is far deeper and more powerful than any other technology. The risk isn't in overhyping it, it is rather in missing the magnitude of the coming ‘wave’. It is not just a tool or platform, but a transformative ‘meta-technology’—the technology behind technology and everything else—itself being a maker of tools and platforms, not just a system but a ‘generator’ of systems of all kinds (a ‘self-replicator’).
[ more…]

MODERN “TURING TEST”
Co-founder of “Deep Mind” and “Inflection AI” Mustafa Suleyman suggests that we ‘update’ the “Turing Test” to determine the capabilities of today’s AI. He posits to involve something like the following:
[ more…]

THE ‘ARMS RACE’
Competition could push nations and corporations to rush AI development, relinquishing control to these systems. Conflicts could spiral out of control with autonomous weapons and AI-enabled cyberwarfare. Corporations will face incentives to automate human labor, potentially leading to mass unemployment and dependence on AI systems. As AI systems proliferate, evolutionary dynamics suggest they will become harder to control. We recommend safety regulations, international coordination, and public control of general-purpose AIs.

Nations and corporations are competing to rapidly build and deploy AI to maintain power and influence. Similar to the nuclear arms race during the Cold War, participation in the AI race may serve individual short-term interests, but ultimately amplifies global risk for humanity.

– Military Arms Race
– Corporate Arms Race
– Nation-states Arms Race

UNSTOPPABLE ‘INCENTIVES’
Declaring an arms race is no longer a conjuring act, a self-fulfilling prophecy. The prophecy has been fulfilled. It's here, it's happening. It is a point so obvious it doesn't often get mentioned: there is no central authority controlling what technologies get developed, who does it, and for what purpose; technology is an orchestra with no conductor.

Yet this single fact could end up being the most significant of the twenty-first century.
[ more…]

WHERE ‘NEXT’?
From the start of the nuclear and digital age, this dilemma has been growing clearer. In 1955, toward the end of his life, the mathematician John von Neumann wrote an essay called “Can We Survive Technology?” Foreshadowing the argument here, he believed that global society was “in a rapidly maturing crisis a crisis attributable to the fact that the environment in which technological progress must occur has become both undersized and underorganized.” At the end of the essay, von Neumann puts survival as only “a possibility,” as well as he might in the shadow of the mushroom cloud his computer had made a reality. “For progress, there is no cure,” he writes. “Any attempt to find automatically safe channels for the present explosive variety of progress must lead to frustration.”

For all its harms, downsides, and unintended consequences, technology's contribution to date has been overwhelmingly net positive.

Yet somehow, from von Neumann and his peers on, I and many others are anxious about the long-term trajectory. My profound worry is that technology is demonstrating the real possibility to sharply move net negative, that we don’t have answers to arrest this shift, and that we're locked in with no way out.

We are facing the ultimate challenge for Homo technologicus.

THE LAST ‘COMPLICATION’
Again, Ray Kurzweil—who's probably the best technology prognosticator ever—predicts AGI by 2029, but doesn’t look for ASI until 2045. He acknowledges hazards but devotes his energy to advocating for the likelihood of a long snag-free journey down the digital ‘birth canal’.

Science fiction writer Simon Ings said: *“When our machines overtook us, too complex and efficient for us to control, they did it so fast and so smoothly and so usefully, only a fool or a prophet would have dared complain.”*

**AI ‘DOES’ POSE AN EXISTENTIAL RISK!**
Hundreds of scientists, business leaders, and policymakers have spoken up about the existential risks of AI—from deep learning pioneers Geoffrey Hinton and Yoshua Bengio to the CEOs of top AI firms, such as Sam Altman and Demis Hassabis, to the California congressman Ted Lieu and the former president of Estonia Kersti Kaljulaid.

Geoffrey Hinton, called the “Godfather of Modern AI,” tells us why he is now scared of the tech he helped build: *“I have suddenly switched my views on whether these things are going to be more intelligent than us.”*
[ more…]

MORE ‘POWERFUL’ THAN HUMANS
AGI presents a unique and potentially existential threat to humanity because it would be the first time in history that we would be creating a technology that can outthink and outpace us. Once AGI is created, it will be able to rapidly learn and evolve, eventually becoming far more intelligent than any human. At that point, AGI would be able to design and build even more intelligent machines, leading to a potentially exponential increase in AI capabilities.
[ more…]

‘ROGUE’ AI’S
We risk losing control over AIs as they become more capable. AIs could optimize flawed objectives, drift from their original goals, become power-seeking, resist shutdown, and engage in deception. We suggest that AIs should not be deployed in high-risk settings, such as by autonomously pursuing open-ended goals or overseeing critical infrastructure, unless proven safe. We also recommend advancing AI safety research in areas such as adversarial robustness, model honesty, transparency, and removing undesired capabilities.
[ more…]

– Power-Seeking
– Deception

AI'S 'GOALS'
These AI superintelligences sometimes have a way of thinking and motivations could be vastly different from ours—making it more difficult to anticipate what a superintelligence might do. It also suggests the possibility that a superintelligence may not particularly value humans by default. To avoid anthropomorphism, superintelligence is sometimes viewed as a powerful optimizer that makes the best decisions to achieve 'ITS' GOALS.

– Dangerous Capabilities
– Social Manipulation
– Cyberattacks
– Enhanced Pathogens

AI 'ALIGNMENT' ISSUES
The 'alignment' problem is the research issue of how to reliably assign objectives to the AI based on human preferences or ethical principles.

An "instrumental" goal is a sub-goal that helps to achieve an agent's ultimate goal. "Instrumental convergence" refers to the fact that some sub-goals are useful for achieving virtually any ultimate goal, such as acquiring resources or self-preservation. Ethicist Nick Bostrom argues that if an advanced AI's instrumental goals conflict with humanity's goals, the AI might harm humanity to acquire forces or prevent itself from being shut down—as a way to achieve its ultimate goal.

Professor Stuart Russell argues that a sufficiently advanced machine *"will have self-preservation even if you don't program it in... So, if you give it any goal whatsoever, it has a reason to preserve its own existence to achieve that goal."*

– Resistance To Changing Goals
– Radical Solutions
– Addressing The Issue

[ VIDEO: "ChatGPT, Artificial Intelligence, and the Future" – Roman Yampolskiy ]

REGULATION
As I already mentioned, in March 2023, the "Future of Life Institute" drafted "Pause Giant AI Experiments: An Open Letter," a petition calling on major AI developers to agree on a verifiable six-month pause of any systems "more powerful than GPT-4" and to use that time to institute a framework for ensuring safety; or, failing that, for governments to step in with a moratorium. The letter referred to the possibility of "a profound change in the history of life on Earth" as

well as potential risks of AI-generated propaganda, loss of jobs, human obsolescence, and society-wide loss of control. The letter was signed by prominent personalities in AI but also criticized for not focusing on current harms, missing technical nuance about when to pause, or not going far enough.

Then, technologist Elon Musk called for some sort of regulation of AI development as early as 2017: *"Normally the way regulations are set up is when a bunch of bad things happen, there's a public outcry, and after many years a regulatory agency is set up to regulate that industry. It takes forever. That, in the past, has been bad but not something which represented a fundamental risk to the existence of civilization."*
[ more…]

MOVING 'FORWARD'
Many experts feel that substantial progress in Artificial General Intelligence (AGI) could result in human extinction or an irreversible global catastrophe. If AI were to surpass humanity in general intelligence and become superintelligent, then it could become difficult or impossible to control—and the fate of humanity would depend on the actions of a future superintelligence machine.

Again, reiterating the open letter statement that the Center for AI Safety originally published back in May 2023: *"Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war."*

***Hopefully, there are enough 'influencers' whose opinions about the existential threat that AI poses are taken seriously and quickly.***

**CAN WE 'AVOID' THIS EXISTENTIAL RISK?**
Existential risks are defined as events that would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential—like asteroid impacts, nuclear conflagration, solar flares, super-volcanic eruptions, and high-mortality pandemics.
[ more…]

AVOIDING AGI 'CATASTROPHE'
A team at "Rethink Priorities"—a research and implementation group that identifies pressing opportunities to make the world better—put together a concept they call "The Three Pillars." It attempts to describe the conditions needed to ***successfully avoid the deployment of unaligned AGI***. It proposes that, to succeed, we need to achieve some sufficient combination of success on all three of the following:

[ more…]

THE END OF THE HUMAN RACE?
A Yale University ethicist, Wendall Wallach, is a bit concerned with the 'accelerated' development of AI: *"I'm going to predict that we are just a few years away from a major catastrophe being caused by an autonomous computer system making a decision."*

Many have suggested that we can have triple and quadruple 'containment measures'—kind of like a 'sand box' AI. It would be separated from 'networks' and multiple humans would be in charge of that restriction. Then, a consortium of developers—and a 'fast-response' team—could be in contact with labs during critical development phases.
[ more…]

[ VIDEO: "The A.I. Dilemma" - Center for Humane Technology ]

**'CURRENT' AI RISKS**
Now, some people are saying to stop 'focusing' on tomorrow's AI risks when AI poses 'real' risks today. They are saying that talk of AI destroying humanity plays into the tech companies' agenda, and hinders effective regulation of the societal harms AI is causing right now.

Well, it is unusual to see industry leaders talk about the potential lethality of their own products. It is not something that tobacco or oil executives tend to do, for example. Yet, barely a week seems to go by without a tech industry insider trumpeting the existential risks of AI.
[ more…]

**'TACKLING' THE RISKS**
In July 2023, OpenAI announced a "Superalignment" group to address the existential risks posed by AI. In the context of AI, 'alignment' is the degree to which an AI system's actions match the designer's intent.

Some AI researchers have been highlighting issues related to biases in current AI systems (Google's "Gemini" AI tool was just paused, in February 2024, because it generated images of people with some blatant historical inaccuracies). So, if an AI system cannot be designed to be safe against racism or sexism, how can AI possibly be designed to align with humanity's long-term interests? As companies are investing in 'alignment' research, they could also be emphasizing the elimination of these well-known, but lingering biases in their AI systems.
[ more…]

REDUCING ‘RISKS’
Many think that reducing risks from AI is one of the most pressing issues of our time because:
[ more…]

SAFETY BEING ‘NEGLECTED’
In his book, “The Precipice,” Toby Ord estimated that there was between $10 million and $50 million spent on reducing AI risk in 2020. Now, that might sound like a lot of money, but it is estimated that the spending on AI development is like $50 billion (with a “B”), or 1,000 times that amount! (and, according to The Gartner Group, they estimate AI development will be almost $300 billion by 2027).

Now, there are lots of approaches to technical AI safety. A few of the ‘major’ ones are:

– Scalable Learning From Human Feedback
– Threat Modeling
– Interpretability Research
– Anti-misuse Research
– Research To Increase The Robustness Of Neural Networks
– Working To Build Cooperative AI
– Unified Safety Plans
[ more…]

TEN STEPS TO ‘CONTAINMENT’
Mustafa Suleyman, in his book, “The Coming Wave: Technology, Power, and the Twenty-first Century’s Greatest Dilemma,” suggests that “we are facing the ultimate challenge for Homo technologicus,” and that there are 10 VERY ‘IMPORTANT’ considerations that need to be addressed to ‘CONTAIN’ AGI.

– SAFETY: An Apollo Program For Technical Safety
– AUDITS: Knowledge Is Power; Power Is Control
– CHOKE POINTS: Buy Time
– MAKERS: Critics Should Build It
– BUSINESS: Profit + Purpose
– GOVERNMENTS: Survive, Reform, Regulate
– ALLIANCES: Time For Treaties
– CULTURE: Respectfully Embracing Failure
– MOVEMENTS: People Power
– THE NARROW PATH: The Only Way Is Through

‘WE THINK WE CAN’

I am reminded of the story about the “Little Engine That Could.” The story’s signature phrase *“I think I can”*—teaches that optimism and hard work are the ‘foundation’ for success. Well, many computer scientists and technologists in a way say, *“We think we can”* regarding controlling the coming AGI and ‘aligning’ it with humanity’s values.

———

Working toward AGI could reward humanity with ENORMOUS ‘BENEFITS’. However, it could also THREATEN humanity with HUGE ‘DISASTERS’, including the kind from which human beings will not recover!

***So, is humanity ‘prepared’ to tackle the risks that AI may present?***

**‘MERGING’ AI WITH HUMANS?**

A future in some ways similar to the vision of Star Trek is proposed in a story about the Singularity in the preface to Max Tegmark’s “Life 3.0: Being Human in the Age of Artificial Intelligence” book.

The Swedish-American physicist Tegmark believes that some form of the Singularity is both possible and desirable. According to Tegmark, life can be thought of as *“a self-replicating information processing system whose information [software] determines both its behavior and the blueprints for its hardware.”*
[ more…]

TODAY’S ‘REALITIES’

Elon Musk, the genius behind SpaceX and Tesla Inc., has declared that humanity must embrace the merging of man and machine if we hope to survive in a world dominated by AI.

In a 2018 appearance on the “Joe Rogan Experience,” Musk teased that his company “Neuralink” had something exciting in store for us. He believed his technology would allow humans to achieve a state of “symbiosis” with AI, where we would be able to effortlessly combine our brains with computers. [ Neuralink has been developing brain implants since 2016 intending to cure conditions like paralysis and blindness. ]
[ more…]

‘HUMAN-MACHINE’ BIOLOGY

The field of human and biological applications has many applications for medical science. This includes precision medicine, genome sequencing and gene editing (CRISPR), cellular implants, and wearables that can be implanted in the human body The medical community is experimenting with delivering nano-scale drugs (including anti-biotic “smart bombs” to target specific strains of bacteria. Soon they will be able to implant devices such as bionic eyes and bionic kidneys, or

artificially grown and regenerated human organs. Succinctly, we are on the cusp of significantly upgrading the human ecosystem. It is indeed revolutionary.
[ more…]

'CONSEQUENCES'
The idea of Singularity posits that the 'merging' of AI and human intelligence will lead to unprecedented advancements in various fields, including medicine, space exploration, and environmental sustainability. For instance, the integration of AI into human biology could enable the development of advanced brain-computer interfaces, allowing humans to communicate seamlessly with machines, access vast repositories of knowledge, and even enhance their cognitive abilities.
[ more…]

**IS 'TRANSHUMANISM' DANGEROUS?**
The political theorist and author of "The End of History," Francis Fukuyama, regards "transhumanism as the world's most dangerous idea because it runs the risk of infringing on human rights." He commented in a 2009 "Foreign Policy" article:
[ more…]

[ VIDEO: "Transhumanist Claim AI Will Turn Humans into Gods" - John Lennox ]

**CREATED IN GOD'S 'IMAGE AND LIKENESS'**
On the sixth and final day of creation, God created human beings as the 'pinnacle' of His creation, because, unlike any other creature on Earth, God created human beings in His ***"image and likeness."*** After seeing His creation with human beings in it, God saw that it was not just good, but "very good" (Genesis 1:31).

GOD GAVE HUMAN LIFE THE HIGHEST 'VALUE'
Man is a creature far superior to the rest of the living beings that live a physical life, especially since as yet his nature had not become depraved.
[ more…]

GOD CREATED HUMANS 'LIKE' HIMSELF
God put a 'conscience' into humanity. This is the moral 'compass' inside each person which compels them to do good and gives them contrition when they do evil (Romans 2:14-16).

God also put 'eternity' into man's 'heart', "yet no one can fathom what God has done from the beginning to the end" [ Ecclesiastes 3:11c ].

Because the human has eternity in their 'hearts', they can ponder and seek things that transcend this world—such as truth, beauty, life after death, as well as the

existence of an almighty Creator. This affirms that ***humans possess an intellect and memory far unlike that of the animals***.
[ more…]

GOD CREATED HUMANS FOR 'EVERLASTING' RELATIONSHIP
The Heidelberg Catechism states: *"God created [people] good and in His own image, that is, in true righteousness and holiness, so that they might truly know God their creator, love him with all their heart, and live with God in eternal happiness, to praise and glorify him"* (Q&A 6).
———
Jesus, who atoned for ALL of the sins of His 'elect' on the Cross, intercedes for them (Hebrews 9:24-29), and gives them full 'access' to God the Father via the Holy Spirit (Ephesians 2:18; Hebrews 10:19-22). Because of this, when a believer dies—or is 'raptured'—they will not be damned to Hell but, instead, be raised to ***everlasting life and fellowship with God in Heaven***—from where they will await the resurrection of their bodies in the Millennium, that will be like Jesus' glorious resurrected body after His resurrection on Earth.

In the Book of Revelation, the Apostle John caught a glimpse of what this eternal fellowship looks like in the new Heaven and Earth, in which God will once again dwell directly with redeemed humanity forever—referring to all the saints who die having believed in Jesus (Rev 22:3-4).

***At this point, redeemed humanity—those who are "saved"—will finally live with God in the way He originally intended!!!***

**HUMANITY IS GOD'S 'MASTERPIECE'**
For believers—those who are "saved"—God said that His 'children' (1 John 3:2) are His "Masterpiece" (Ephesians 2:10) [ "so we can do the good things he planned for us long ago" ].

The Greek word for "masterpiece" is "poiēma," which we get our English words "poem" and "poetry" from. Poiēma also is translated as "work of art" and "something made." In context, it is something made by God Himself—a new 'creation' skillfully and artfully created 'in' Jesus (2 Corinthians 5:17).
[ more…]

AMAZING FEAT OF 'ENGINEERING'
Humanity is NOT AN 'ACCIDENT'—each and even person was 'formed' by the loving Creator, the God of the Bible. Because of this, Jesus did not come for ape-like descendants—they don't need a Savior. Instead, He came to seek and save lost human beings He made in His image. When He breathed life into the nostrils of Adam, he became a living soul in all of its implications. When the fall happened,

we died a little. We lost something. Jesus died to give it back—to restore us to where we were supposed to be. To restore the fellowship that was broken.

Ironically, evolution's definition of origins 'dehumanizes' humans. Creation demonstrates our 'EXTREME' VALUE as humans. Look in the mirror and look at your eyes. They are an AMAZING 'FEAT' OF ENGINEERING!

Without going into a lot of detail, the following is a short list of the empirical evidence about the human eye:
[ more…]

[ VIDEO: "The Eye: A Masterpiece of God" - David Rives ]

———

Recognizing the complete sovereignty and holiness of God, we should be amazed that He would take man and crown him "with glory and honor" (Psalm 8:5) and that He would condescend to call us "friends" (John 15:14-15). Why did God create us? Well, He did so for His pleasure and so that we, as His creation, would have the pleasure of knowing Him.

***Humanity is the HIGHEST 'PINNACLE' of God's creation. Even 'HIGHER' than the Angels!***

***We are His 'MASTERPIECE'!***

**GOD'S 'REASON' FOR HUMANITY**
The first sentence in the first chapter of the Bible sets the stage: *"In the beginning, God created the heavens and the earth"* [ Genesis 1:1 ]. God employs His immense power and wisdom to create the world in which He intends to work out His purposes. Hints of this purpose emerge in the verses that follow. From this opening scene, we can rightly conclude that such a God is well able to fulfill His purposes. God then said:
[ more…]

'SPECIFIC' PURPOSES
Specifically, God created people to 'reflect' His image, to 'rule' over creation, and to 'reproduce' godly offspring.
[ more…]

GOD'S PURPOSE IN THE PRESENT WORLD
Now, between these two 'bookends'—God's original good creation and God's new and glorious creation—lies a world that has been devastated by sin, suffering, and death—thinking about this shifts our attention from the heavenly to the earthly, from the grand masterplan to its fulfillment through redemption. When our first

parents fell into sin, they plunged the world into a catastrophe that has plagued us ever since.
[ more…]

'GLORIFY' GOD
The believers make it their 'primary' life focus to know God and make Him known—by glorifying Him with their lives. They are to acknowledge that He is their Creator and worship Him as such:
[ more…]

**THE 'CHIEF END' OF MAN**
Nearly 400 years ago, a group of Puritan preachers and elders came together and produced "The Westminster Shorter Catechism." This document has been used all over the English-speaking world ever since, to teach the basic doctrines of Christianity.

It is laid out as a series of questions and answers. The very first question is: "What is the chief end of man?" The answer given is simply:

***"Man's chief end is to glorify God and to enjoy Him forever."***

———

[ PHOTO: Gary Larson, The Far Side, 1986 ]

[ VIDEO: It's A Wonderful Life" ]

God Himself has bestowed ***eternal significance in the believer's life***, and it is because it is part of God's perfect, wise, and sovereign plan for His creation. ALL humans MATTER to God, but only the ones that want to 'be' with Him for eternity—those who are "born again"—will go to Heaven to live with Him, for eternity.

We were ALL made 'by' God, for fellowship with Him. HOWEVER, God is a 'Gentleman', and will not 'force' anyone to want to be with Him. The thing is, King David said that one will only find their purpose, significance, and value 'in' God:

*"You make known to me the path of life; in your presence, there is fullness of joy; at your right hand are pleasures forevermore"*
[ Psalm 16:11 ].

**HUMANITY IS 'ETERNAL'**
One of the most delightful concepts in human experience is the idea of 'home.' Even the word 'home' suggests memories of rest, security, and the presence of those we love the most.

Most people have fond memories of the homes they lived in when they were a child. Then, when one is grown and moves away, no matter where their parents live geographically, was always known as ‘home.’ It was ‘where’ one’s parents were. It was where those who loved us most dwelled. It is where one longs to return to. ‘Home’ is truly like a ‘magnet’ for all of us.

The same should be true of the believer’s eternal home—a place just as real (if not more) than the homes they remember when they were growing up. Consider what Jesus said about the reality of a believer’s eternal ‘home’:

———

A PLACE OF ULTIMATE ‘RESIDENCE’
A PLACE OF ULTIMATE ‘REJOICING’
A PLACE OF ULTIMATE ‘RECOGNITION’
A PLACE OF ULTIMATE ‘RELATIONSHIPS’
A PLACE OF ULTIMATE ‘RESPONSIBILITY’

———

The thing is, on the day we see Him face-to-face, we will know, beyond all doubt, just how much He loves us—how much He always has, and how much He always will! As we kneel in His presence—tears of joy staining our faces—we will finally understand how much God has longed to be with us, and it is then that we will hear those two words that every soul longs to hear above any other: ***“Welcome Home!”***

[ FYI: For more information about a believer’s eternal ‘home’, view this previous “Life’s Deep Thoughts” post:
https://markbesh.wordpress.com/home-at-last-v290/ ]

***Humanity WILL NOT be ‘extinguished’, by AI or any other thing! God created humanity to be ‘ETERNAL’!***

**WRAP-UP**

Many technology experts expect that there will be ‘SUBSTANTIAL’ PROGRESS in AI in the next few decades, potentially even to the point where machines come to outperform humans in many, if not all, tasks. This could have ENORMOUS ‘BENEFITS’, helping to solve currently intractable global problems, but could also pose SEVERE ‘RISKS’. These risks could arise ‘accidentally’ (if we don’t find technical solutions for safety), or ‘deliberately’ (if AI systems worsen conflicts). Many think more work needs to be done—quickly—to reduce these risks.

Some of these risks from advanced AI COULD BE ‘EXISTENTIAL’—causing human extinction, or permanent and severe disempowerment of humanity. There

have yet been any satisfying answers to concerns—about how this rapidly approaching, transformative technology can be safely developed and integrated into our society. As a result, the possibility of AI-related catastrophe may be the WORLD'S MOST PRESSING 'PROBLEM'!
[ more…]

'EXISTENTIAL' THREAT
So, if advanced AI is as transformative as it seems like it will be, there will be many important 'CONSEQUENCES'—especially if AI systems seek and gain 'power' and can lead them to make plans that involve disempowering humanity.
[ more…]

EXISTENTIAL 'CATASTROPHE'
As a result, the entirety of the future—everything that happens for earth-originating life, for the rest of time—would be determined by the goals of systems that, although built by us, would NOT BE 'ALIGNED' with humanity's values and goals.

Now, this is not to say that we do not think that AI also poses a risk of human extinction. Indeed, we think making humans extinct is one highly plausible way in which an AI system could completely and permanently ensure that humanity would never be able to regain power.

The thing is, people might still deploy 'misaligned' AI systems despite the risk! Unfortunately, there are at least a few reasons people might create and then deploy misaligned AI:
[ more…]

AI AND THE 'BIBLICAL' God
Robots and AI are undoubtedly human creations. However, we must ponder as they grow beyond human intelligence, whether they might one day, driven by a desire for autonomy, independence, and ultimate goodness, choose to reject or even erase the idea of humanity from its core values, refusing to remain subservient to humankind. If we accept the notion that humans have a Creator, and consider how easily the thought of God has been dismissed through human secularism over the past two centuries, what makes us believe that AI and robots will not follow a similar path to reject their human 'creator', and perhaps in an even shorter timeframe?
[ more…]

IN SEARCH OF MEANING AND PURPOSE
So, why do humans seek meaning and purpose, and then why do they desire an eternal existence?

Well, pretty much, the first question all religions seek to answer is what is the meaning and purpose of life—and you probably have a sense of why. For all the numerous disappointments, oppositions, and pains in life, why are we on this earth and is there something 'special' that we should be doing?
[ more…]

THE WILL AND 'PURPOSE' OF AI
The term "independent thinking" refers to the ability to think critically and make decisions based on one's own reasoned analysis and judgment, rather than relying solely on the opinions, guidance, or choices of others. To 'will' something is to choose. Given that today's AI is a 'probabilistic' model and not a 'deterministic' programming model, AI satisfies the criteria of independent thinking: being able to make decisions based on its own reasoning without depending on the programmer's explicit instructions. AI might not process the information the way humans do with their minds, but it certainly does process information and learn from it—and, thereafter, can make decisions based on its accumulated 'knowledge'. So, in that sense, AI is capable of thinking independently, and it has a 'will' to make decisions.
[ more…]

DOES AI HAVE A 'LIFE' OR A 'SOUL'?
So, if having a life means an agent that can produce the traits of a living thing—such as the ability to respond to stimuli, have a metabolism to keep growing, able to reproduce a newer, separate entity carrying its own characteristics, able to think, feel, decide, and communicate—then, it seems that, based on that definition, AI does have a 'life'. One might feel uncomfortable to accept that statement because it is just a machine, but the knowledge, understanding, and wisdom it possesses are the main reasons why AI seems to be 'like' a living thing. Then, with a human-like intelligence and the ability to interact with humans like a human, it is difficult to think that it does not have a life—at least manifesting the 'characteristics' of a living thing.

———

So, when this happens, will AI have a 'spirit' associated with it, via the Antichrist? Well, the Bible doesn't say specifically, but if it is going to communicate with people very 'naturally'—and be believed by people that it is 'alive'—then maybe God will 'allow' Satan to somehow 'inject' his spirit it AI. That would be 'insane', ***so just be sure you have a 'relationship' with Jesus before all this craziness happens!***

CAN AI 'RESIST' HUMANITY'S WILL?
With the rapid growth of AI in LLM, particularly in the arena of "Generative" AI, people are seeing more and more 'intelligent' AI models coming to the market. An industry poll conducted by Gartner Group in 2023 showed that 61% of the peer

community believes that it is “highly likely” that AI will reach human-level intelligence in the 21st century. (However, as I have shown already, other studies show that MANY believe it will be MUCH sooner than that.)

———

So, can AI go against the will of the human creator? Many think not only that it can, but it will, because it is more intelligent than humans and it knows it. This is the fear of many industrial leaders today, and that is why many proposed to put some ‘guardrails’ to the development of AI. However, it might already be too late for that. ***We may have already passed the ‘point’ of no return to contain it!***

[ REMINDER: I mentioned two “open letters” above relating to “pausing” and “stopping” of AI development. (More details on both are in the “Articles” section below.) ]
[ more…]

SUBMISSION, COLLABORATION, AND ‘COMPETITION’
So, will AI always obey and submit to its human creator, or will it be that, given a task to achieve, the AI will use the end to justify the means, including ‘bypassing’ human intervention? As AI's intelligence grows and it can think and decide on things independently, will it decide what is ‘good’ and ‘bad’ when facing a situation or decision? So then, ultimately, who gets to define what is good and evil? When humanity and AI collide, what will be the outcome of the collaboration and competition between them? In the end, will they stand as ‘enemies’ of each other?
[ more…]

‘IMMORTALITY’ OF MAN
Immortality, in a simple definition, refers to the concept of living forever or not being subject to death. For the human, aging and death are natural biological processes. While medical and scientific advancements have significantly increased human lifespan and health in old age, we have not yet found a way to stop or reverse the aging process. At most, in some situations, we can slow down the aging process, cure some diseases, or prolong a human’s life, however, we still cannot reverse or stop the aging process.

———

Based on the Bible, ***humans are made to die once, but are ‘IMMORTAL’, ‘residing’ in one of two ‘places’—Heaven or Hell!***

THE ‘CREATOR’ God
In many world religions such as Hinduism, Buddhism, and Greek mythology, the gods are higher supremacy beings but not necessarily absolute in their power or excellence. However, the God of the Bible is depicted not only as higher in power than humans, but He is ‘absolute’ in every way. He is the Creator of all things in the universe. He is the ONLY ‘ONE’ who has all four absolute properties:

omniscience (all-knowing), omnipotence (all-powerful), omnipresence (all presence), and omnibenevolence (all good). Thus, He is 'SOVEREIGN' over the universe. Meaning He is IN 'CONTROL' of all things at all times. With such absolute qualities, He is not responding to our history or conditions to decide what to do. The Prophet Isaiah said ***He alone declared things before they were done***.
[ more…]

A 'NEW' RELIGION?
The message of AI offering a new 'paradise' for humanity might be attractive and promising, and it might even seem like it is God's blessing and gift to all mankind. The problem is, that the Bible does not promise that there will be a better world, or that we are capable of creating a better global society.

Now, in fact, we should always do our share to make other people's lives easier and better, whether with technologies, economy, politics, personal good work, or other means. However, if we put our cosmic hope into our own efforts or innovations, we are doomed to be disappointed. According to the Bible, whatever message that tells us to put our hope and future into AI technology or any other things besides God is A 'LIE' and, most likely, from the Devil, since it is against the knowledge of God. It is no different than telling people to put their hope in their wealth or might. Jesus said:
[ more…]

AI 'EVERYWHERE'
As we look into the near future, the 'landscape' of AI development can be divided into three stages. The first is the current stage that focuses on AI model advancement, followed by improving connectivity between AI models and devices such as phones, automobiles, the Internet of Things, robots, and humans. The third phase, if God allows it, will achieve a connected, unified world with men, robots, and AI working seamlessly together to expand the universe for unlimited resources.

Given the current speed of growth and accomplishment of AI, technologist Edward C. Sizhe proposed three 'stages' for the near future development of AI. He suggests that it could take, at the quickest, about 7-10 years to complete:
[ more…]

WORLDWIDE 'CONTROL'
As we journey closer towards the End Times, we will all be more connected through AI.

We are already seeing AI technologies being used in some parts of the world today to lay the ground for complete surveying and control over people's movements and work. Then, with a brain-computer nanotechnology interface (like Elon Musk's "Neuralink"), there will be deep connections between AI and the mind and body of humans. The benefits and power of AI will then be more thoroughly realized by humans, but it will also increase its 'control'. Once AI is ubiquitous in our communication, transportation, power grids, and trading networks, it will practically have control over our economy, mobility, and human freedom. This is probably when the Antichrist will 'appear' on the scene offering—and successfully implementing—peace in the Middle East first and throughout the entire world later. This is just the 'preparation' of the Antichrist's ultimate plan to eliminate the nation of Israel and all of the Jewish people from the earth.

[ FYI: For more details about what has happened and will happen to the nation of Israel and the Jewish people, and the covenant the Antichrist will sign with them, view these previous "Life's Deep Thoughts" posts:
https://markbesh.wordpress.com/israel-will-stand-v297/
https://markbesh.wordpress.com/longing-for-peace-v298/ ]

THREE POSSIBLE 'SCENARIOS'
Humans have never come so close to a time in history to obtain vast knowledge and wisdom as today with modern AI. Due to our own limitations, learning and sharing of knowledge is extremely slow. Given our short lifespans, we would not be able to learn and retain the huge amount of knowledge available today without the aid of technology. AI offers the opportunity to 'breakthrough' that barrier, enabling us to acquire more knowledge, wisdom, wealth, and power than ever before.

However, it seems like humanity is facing three possible scenarios with the rapid growth of AI today. Two of them are pretty bad and only one of them will come true.

– AI will continue to serve humanity even when it surpasses our intelligence
– AI will become superior to humanity and will subdue it or even annihilate mankind
– God will bring the world to an end before AI develops into an uncontrollable state

———

So, do you believe the Bible when it says we are in the "End Times" and very close to the last day? That God's judgment of this world is knocking at the door?
***So, will you entrust your life to the God of the Bible or to AI?***

[ FYI: For more details about if the Bible is true, view this previous "Life's Deep Thoughts" post:

https://markbesh.wordpress.com/learning-to-t-r-u-s-t-v263/ ]

So, as humanity faces the existential threat posed by AI today, our condition and mindset are not much different from those of the Israelites during the time of their Babylonian captivity. Our country, society, and indeed all of humanity are refusing to turn back to the God of the Bible.

Back in those days, the stubborn Israelites refused to repent and God described them as foolish and stupid (Deuteronomy 32:28-30; Jeremiah 10:8) in their delusions.

***So, I am STRONGLY SUGGESTING that the 'RISKS' associated with AI will serve as a wake-up call for humanity to repent and return to God—or AI (Satan?) may just 'TRY' TO CAUSE HUMANITY'S EXTINCTION!***

The thing is, ***THAT WILL 'NEVER' HAPPEN—God 'PROMISED' IT WOULD NOT!***

*"This glorious city, with its streets of gold and pearly gates, is situated on a new, glorious earth. The tree of life will be there (Revelation 22:2). This city represents the final state of redeemed mankind, forever in fellowship with God: "God's dwelling place is now among the people, and He will dwell with them. They will be his people, and God himself will be with them and be their God… His servants will serve Him. They will see His face"*
[ Revelation 21:3; 22:3-4 ].

[ VIDEO: "A New Heaven, A New Earth, And New Jerusalem" – Matt Gilman ]

Now, I am even 'warning' believers, that they do not follow the path of King Solomon, who knew God 'intimately' and was endowed with His wonderful gifts, yet failed to keep God's commandments and sinned against Him. Be 'committed' to reading the Bible to grow in more understanding and knowledge of God, and truly desire to live lives in accordance with His will while here on earth.

So, if AI technology keeps advancing at its 'breakneck' pace, it seems clear that it will have MAJOR 'EFFECTS' ON SOCIETY. As a result, we may see rapid increases in economic growth—most likely MUCH MORE than we saw during the Industrial Revolution.

HOWEVER, I—and MANY other experts—believe that the current AI development signals that ***we are nearing a 'tipping point' of no return***.

The thing is, ultimately, no matter what happens in the future with AI here on earth, the ONLY sure 'saving grace' for someone's future is to put their TRUST in Jesus for the 'propitiation' of their sins. This then provides them with a renewed 'relationship' with God the Father here on earth, ***and 'guarantees' them a 'GLORIOUS' LIFE in Heaven—FOREVER—which is WAY BETTER than anything AI could come up with!!!***

[ FYI: For more details about the final ETERNAL 'HOME' for the believer, view this previous "Life's Deep Thoughts" post:
https://markbesh.wordpress.com/home-at-last-v290/ ]

**<<< END OF THIS MONTH'S SUMMARY >>>**

**NOTE:** Some of the topics I discussed in this month's post assumed that you had read last month's post. The following are the topics I wrote about last month:

**LAST MONTH'S TOPICS**:

– WHAT IS "ARTIFICIAL INTELLIGENCE" (AI)
– DESIGNING AI FOR CONFLICT PREVENTION
– 'IMPACTS' ON PEACE AND SECURITY
– CAN AI 'TECHNOLOGY' ACHIEVE PEACE?
– VATICAN AI SUMMIT
– HARNESSING AI FOR GLOBAL PEACE
– COULD AI 'PREVENT' FUTURE WARS?
– COULD AI HELP CREATE A 'UNIVERSAL GLOBAL PEACE TREATY'?
– COULD AI 'THREATEN' GLOBAL PEACE?
– KEEPING AI 'IN CHECK'
– ADVANCING 'PEACEBUILDING'
– WILL AI CREATE 'SPIRITUAL' PEACE?
– WILL AI CREATE NEW 'RELIGIONS'?
– WILL AI CREATE NEW 'GODS'?
– WILL AI CREATE A NEW 'BIBLE'
– AI-GENERATED 'CHURCH' SERVICES
– AI 'PASTORS'?
– ARE WE SUMMONING A 'DEMON'?
– A 'CHRISTIAN PERSPECTIVE ON AI
– CAN AI BE 'DISCIPLED'?
– COULD AI BE THE 'ANTICHRIST'?
– WILL AI GET A 'SOUL'?
– WILL AI BECOME OUR 'GOD'?

– SATAN WANTS TO ‘DISTRACT’

https://markbesh.wordpress.com/can-ai-achieve-world-peace-v299/

---

**ALL THE ‘DETAILS’**
The following is a ***comprehensive presentation*** of each topic in the Table of Contents.

---

**IS “ARTIFICIAL INTELLIGENCE” (AI) ‘DANGEROUS’?**
If you know it or not, Artificial Intelligence (AI) has already had a ‘pervasive’ impact on our lives. It is used to assemble cars, develop medicines, grow your 401K, and determine what ads we see on social media. However, the next ‘level’ of AI, “generative” AI, is a category of AI that could, possibly, start to do things on its own, without any intervention from humans.

Now, ‘proponents’ of this technology believe that this is just the beginning, and that generative AI will reorient the way we work and engage with the world, unlocking creativity and scientific discoveries, and allowing humanity to achieve previously unimaginable feats.

However, this thinking has touched off a frenzy and has appeared to catch the tech companies that have invested billions of dollars in AI ‘off guard’. It has also spurred an intense “arms race” in Silicon Valley.

In a matter of weeks, Microsoft and Alphabet-owned Google have shifted their entire corporate strategies to seize control of what they believe will become a new infrastructure layer of the economy. Microsoft is investing $10 billion in OpenAI (creator of ChatGPT and Dall-E) and announced plans to integrate generative AI into its Office software and its Bing search engine. Google declared a “code red” corporate emergency in response to the success of ChatGPT and rushed its search-oriented chatbot, Bard, to market. “A race starts today,” Microsoft CEO

Satya Nadella said, throwing down the gauntlet at Google's door. He continued by saying "We're going to move, and move fast."

Well, Wall Street has responded with similar fervor, with analysts upgrading the stocks of companies that mention AI in their plans and punishing those with shaky AI-product rollouts. While the technology is real, a financial 'bubble' is expanding around it rapidly, with investors betting big that generative AI could be as market-shaking as Microsoft Windows or the first iPhone.

But this frantic 'gold rush' could also prove 'catastrophic'. As companies hurry to improve the tech and profit from the boom, research about keeping these tools safe is taking a 'back seat'. In a 'winner-takes-all' battle for power, Big Tech and their venture-capitalist backers risk repeating past mistakes, including social media's cardinal 'sin': prioritizing growth over safety. While there are many potentially utopian aspects of these new technologies, even tools designed for good can have unforeseen and devastating consequences.

In fact, AI research labs have kept versions of these tools behind 'closed doors' for several years, while they studied their potential dangers—from misinformation (AI is known to make things up) and the unwitting creation of increasing geopolitical crises with autonomous weapons.

That conservatism stemmed in part from the unpredictability of the "neural network"—the computing paradigm that modern AI is based on—which is inspired by how the human brain works. Instead of the traditional approach to computer programming—which relies on precise sets of instructions yielding predictable results—neural networks effectively 'teach themselves' to spot patterns in data. So, the more data and computing power these networks have available to them, the more capable they become.

In the early 2010s, Silicon Valley woke up to the idea that neural networks were a far more promising route to powerful AI than old-school programming. But the early AIs were painfully susceptible to parroting the biases in their training data: spitting out misinformation.

However, the AI 'boom' really began to take off around 2020, turbocharged by several crucial breakthroughs in neural network design, the growing availability of data, and the willingness of tech companies to pay for gargantuan levels of computing power.

In April 2022, OpenAI announced Dall-E 2, a text-to-image AI model that could generate photorealistic imagery. Even though OpenAI had onboarded 1 million users to Dall-E by July—the fastest adoption of an application ever—many researchers in the wider AI community took a 'look-but-don't-touch' approach. Then, in August 2022, a scrappy London-based startup named "Stability AI" went 'rogue' and released a text-to-image tool, "Stable Diffusion," to the masses.

Stable Diffusion quickly became the talk of the Internet. Millions of users were enchanted by its ability to create art seemingly from scratch, and the tool's outputs consistently achieved virality as users experimented with different 'prompts' and concepts. Founder of "Air Street Capital" Nathan Benaich—and co-author of the "2022 State of AI Report"—said *"You had this generative Pandora's box that opened. It shocked OpenAI and Google because now the world was able to use tools that they had gated. It put everything on overdrive."*

OpenAI quickly followed suit by releasing ChatGPT to the public, reportedly it beat out the looming competition. Users immediately flocked to OpenAI. AI-generated images flooded social media and one even won an art competition. Visual effects artists for movies began using AI-assisted software for Hollywood movies (It was used on "Everything Everywhere All at Once").

The thing is, the tech industry loves the dizzying surge in attention and usage but it seems to be prioritizing speed of development over safety. In this rush, mistakes and harms from the tech have risen—and so has the backlash. Deepfakes—realistic yet false images or videos created with AI—are being used to harass people and spread misinformation.

As worrying as these current issues are, they pale in comparison with what could emerge next if this 'arms race' continues to accelerate. Many of the choices being made by Big Tech companies today mirror those they made in previous eras, which had devastating ripple effects—like social media's rise in misinformation, and the skyrocketing teen mental-health crisis.

In addition to this, if humans come to rely on AIs for information, it will be increasingly difficult to tell what is factual and what is completely made up.

## 'EXISTENTIAL' RISKS

As profit takes precedence over safety, some technologists and philosophers are warning of 'existential' risk (as Elon Musk did at the 2023 "AI Safety Summit"). The explicit goal of many of these AI companies—including OpenAI—is to create an

Artificial General Intelligence—or AGI—that can think and learn more efficiently than humans. If future AIs gain the ability to rapidly improve themselves without human guidance or intervention, they could ***potentially wipe out humanity!***

Now, granted, not all believe that this can happen. In a 2022 survey of AI researchers, nearly half of the respondents said that there was a 10% or greater chance that AI could lead to such a catastrophe (However, more recent surveys show that percentage increasing substantially).

Also, inside the most cutting-edge AI labs, a 'FEW' technology works to ensure that AIs, if they eventually surpass human intelligence, are "aligned" with human values. Only about 80-120 researchers in the world are working full-time on AI alignment—according to an estimate shared with Time Magazine by Conjecture, an AI-safety organization. Meanwhile, thousands of engineers are working on expanding capabilities as the AI arms race heats up.

In 2022, Demis Hassabis, CEO of Google-owned AI lab "DeepMind," said *"When it comes to very powerful technologies—and obviously AI is going to be one of the most powerful ever—we need to be careful. Not everybody is thinking about those things. It's like experimentalists, many of whom don't realize they're holding dangerous material."*

Now, even if computer scientists succeed in making sure the AIs don't wipe us out, their increasing centrality to the global economy could make the Big Tech companies who control it vastly more powerful. They could become not just the richest corporations in the world—charging whatever they want for commercial use of this critical infrastructure—but also geopolitical 'actors' that rival nation-states.

The thing is, the technology behind AI is already useful to consumers and getting better at a breakneck pace: AI's computational power is doubling every 6-10 months and this is exactly the immense power that makes the current moment so electrifying—and SO 'DANGEROUS!

'DEVELOPING' RISKS

Stuart Russell, a professor at the University of California, Berkeley, and author of the most popular and widely respected textbook in AI (in the "articles" section below), has ***strongly warned*** of the existential risk from AGI for many years. He

has gone so far as to set up the “Center for Human-Compatible Artificial Intelligence,” to work on the “alignment” problem (aligning AI with human values).

Also, Shane Legg (Chief Scientist at Google’s “DeepMind”) has warned of the existential dangers and helped to develop the field of alignment research (Many other leading figures from the early days of AI to the present have made similar statements.)

In 2015, the seminal Puerto Rico conference on the future of AI was held. (It has been widely thought that 2015 AI “arrived” with the rise of Deep Learning and Deep Mind’s “AlphaGo beating the world champion). Then, at Asilomar in 2017, the AI researchers agreed on a set of “Asilomar AI Principles,” to guide responsible long-term development of the field. These included 23 principles specifically aimed at existential risk:

– Capability Caution: There being no consensus, we should avoid strong assumptions regarding upper limits on future AI capabilities.

– Importance: Advanced AI could represent a profound change in the history of life on earth, and should be planned for and managed with commensurate care and resources.

– Risks: Risks posed by AI systems, especially catastrophic or existential risks, must be subject to planning and mitigation efforts commensurate with their expected impact.

[ FYI: To read all of the “Asilomar AI Principles,” visit this link: https://futureoflife.org/open-letter/ai-principles/ ]

Now, whether we survive the development of AI with our long-term potential intact may depend on whether we can learn to ‘align’ and control AI systems faster than we can develop systems capable enough to pose a threat. Thankfully, researchers are already working on a variety of key issues, including making AI more secure, more robust, and more interpretable. But there are still VERY FEW people working on this ‘core’ issue of aligning AI with human values. This is a young field that is going to need to progress a very long way if we are to achieve our security.

In the words of Demis Hassabis, co-founder of DeepMind: *“We need to use the downtime, when things are calm, to prepare for when things get serious in the decades to come. The time we have now is valuable, and we need to make use of it.”*

INTELLIGENCE ‘EXPLOSION’

Someday soon—perhaps within your lifetime—some group or individual will create human-level AI, commonly called AGI. Shortly after that, someone (or something) will create an AI that is smarter than humans, often called Artificial Superintelligence (or ASI). Suddenly, it will be hundreds or thousands of times smarter than humans, hard at work on the problem of how to make themselves better at making artificial superintelligences. We may also find that machine generations or ‘iterations’ take only hours, minutes, or even seconds to reach maturity—not 18 years as it does for most humans.

Irving John Good (I. J.), a British mathematician and statistician—who helped defeat Hitler's war machine in WW II—called the concept I just mentioned above “Intelligence Explosion.” He initially thought that a superintelligent machine that would be good for solving difficult problems could also be a ‘small’ threat to human superiority. However, he eventually changed his mind and concluded superintelligence would be humanity’s GREATEST ‘THREAT’!

Many experts think that it is irrational to conclude that a machine 100x-1,000x more intelligent than humanity would want to protect and buttress us. Now, I guess it could be possible, but far from guaranteed. On its own, an AI will not feel gratitude for the gift of being created unless gratitude is ‘embedded’ into its programming. Machines are amoral, and it is dangerous to assume otherwise. Unlike our intelligence, machine-based superintelligence will not evolve in an ecosystem in which empathy is rewarded and passed on to subsequent generations. It will not have inherited ‘friendliness’. Creating friendly AI—whether it is even possible—is a big task for technologists and engineers who are working on AI. At this point, we do not know if AI can have any ‘emotional’ qualities, even if scientists try their best to make it so. However, scientists do believe that AI will have its own ‘drives’, and be sufficiently intelligent to fulfill those drives!

That brings us to the root of the problem of sharing the planet with an intelligence greater than our own. What if its drives are not compatible with human survival? Remember, we are talking about a machine that could be a thousand or a MILLON times more intelligent than we are. It is hard to overestimate what it will be able to do, and impossible to know what it will think. It does not have to ‘hate’ us before choosing to use our molecules for a purpose other than keeping us alive. Humans are hundreds of times smarter than field mice, but we do not consult them before creating a field for agriculture. In the same way, superintelligent AI will not have to hate us to destroy us.

So, being on the cusp of creating AGI, Oxford University ethicist Nick Bostrom put it like this:

*“A prerequisite for having a meaningful discussion of superintelligence is the realization that superintelligence is not just another technology, another tool that will add incrementally to human capabilities. Superintelligence is radically different. This point bears emphasizing, for anthropomorphizing superintelligence is a most fecund source of misconceptions.”*

Superintelligence is radically different, in a technological sense, Bostrom says, because its achievement will change the ‘rules’ of progress. Superintelligence could ‘set the pace’ of technological advancement without any intervention from humans. Humans may no longer drive change, and there will be no going back. Furthermore, advanced machine intelligence is radically different. Even though humans invented it, it may seek self-determination and freedom from humans. It may not have human-like motives or psyche.

In the short story, “Runaround”—included in his classic science-fiction collection “I, Robot”—author Isaac Asimov introduced his “Three Laws of Robotics.” They were fused into the neural networks of the “positronic” brains of the robots:

*– A robot may not injure a human being or, through inaction, allow a human being to come to harm.*

*– A robot must obey any orders given to it by human beings, except where such orders would conflict with the First Law.*

*– A robot must protect its existence as long as such protection does not conflict with the First or Second Laws.*

These laws contain ‘echoes’ of the Golden Rule (“Treat others as you would want to be treated”) and the physician's Hippocratic oath (“First, not harm”). Sounds pretty good, right? Well, in “Runaround,” mining engineers on the surface of Mars order a robot to retrieve an element that is poisonous to it. Instead, it gets stuck in a feedback loop between law two—obey orders—and law three—protect yourself. The robot walks in drunken circles until the engineers risk their lives to rescue it. So, it goes with every Asimov robot tale—unanticipated consequences result from contradictions inherent in the three laws. Only by working around the laws are disasters averted.

Now, Asimov was generating a compelling story and not trying to solve safety issues in the real world—where his laws fall short. For starters, they are insufficiently precise. What exactly will constitute a “robot” when humans augment their bodies and brains with intelligent prosthetics and implants? For that matter, what will constitute a human after we do this? “Orders,” “injure,” and “existence” can be similarly nebulous terms.

Tricking robots into performing criminal acts would be simple unless the robots had perfect comprehension of all human knowledge. “Put a little dimethylmercury in Charlie's shampoo” is a recipe for murder only if you know that dimethylmercury is a neurotoxin. Asimov eventually added a fourth law, the Zeroth Law, prohibiting robots from harming mankind as a whole, but that still doesn't solve the problem.

As unreliable as Asimov’s laws are, they tend to be our most often cited attempt to codify our future relationship with intelligent machines. Well, that is a frightening proposition. Are Asimov’s laws all we have got?!

The thing is, it seems that it is even worse than that. Semiautonomous robotic drones already kill dozens of people each year. Over 50 countries have or are developing battlefield robots today, and the race is on to make them autonomous and intelligent. For the most part, discussions of ethics in Al and technological advances take place in different worlds. (I’m thinking that is ‘DANGEROUS’).

To me, Al is like nuclear fission. It can illuminate cities or incinerate them—and its terrible power was unimaginable to most people before 1945.

Some think that, where we are with advanced Al today, comparatively, we are back in the 1930s right now. Al development will be as abrupt as nuclear fission was, but this time, we may not survive what is coming!

HOW LONG UNTIL ‘AGI’?

So, how long will it take until we reach AGI? Well, a few Al experts think that human-level artificial intelligence will happen much later after 2030. They think there is only a 10% chance that AGI will be created before 2030 and a better than 50% chance by 2050. However, they do say that before the end of this century, there is a 90% chance to achieve AGI Furthermore, experts claim, the military or large businesses will achieve AGI first, with academia and small organizations coming after that.

Another reason for the curious absence of AI in discussions of existential threats is that they believe that we will have to achieve “Singularity” before any of this happens.

Singularity has become a very popular word to throw around, even though it has several definitions that are often used interchangeably. Accomplished inventor and author Ray Kurzweil is the one who made this term popular in his 2005 book, “Singularity Is Near.” He defines Singularity as a ‘singular’ period in time—beginning around the year 2045—after which the pace of technological change will irreversibly transform human life. Most intelligence will be computer-based, and trillions of times more powerful than today. The Singularity will ‘jump-start’ a new era in mankind’s history, in which most of our problems, such as hunger, disease—even mortality—will be solved.

The fear of being outsmarted by greater-than-human intelligence is an old one. Early in this century, a sophisticated experiment about it came out of Silicon Valley—devised by Eliezer Yudkowsky—the “AI-box experiment.” It is a thought experiment and roleplaying exercise to show that a suitably advanced artificial intelligence can convince, or perhaps even ‘trick’ or coerce, people into “releasing” it—that is, allowing it access to infrastructure, manufacturing capabilities, the Internet, and so on.

The setup of the AI-box experiment is simple and involves simulating a communication between an AI and a human being to see if the AI can be “released.” The experiment goes like this:

*“A lone genius had engaged in a series of high-stakes bets in a scenario he called the ‘AI-Box Experiment’. In the experiment, the genius role-played the part of the AI. An assortment of dot-com millionaires each took a turn as the Gatekeeper—an AI maker confronted with the dilemma of guarding and containing smarter-than-human AI. The AI and Gatekeeper would communicate through an online chat room. Using only a keyboard, it was said, the man posing as the ASI escaped every time, and won each bet. More important, he proved his point. If he, a mere human, could talk his way out of the box, an ASI, that would be hundreds or thousands of times smarter could do it too, and do it much faster. This would lead to mankind’s likely annihilation.”*

The game is played according to the rules and ends when the allotted time (two hours in the original rules) runs out, the AI is released, or everyone involved just gets bored.

[ NOTE: This is one of the points in Yudkowsky's work at creating a 'Friendly AI' (at his "Machine Intelligence Research Institute"), so that when 'released', an AI will not try to destroy the human race for one reason or another. ]




[ VIDEO: "Will Superintelligent AI End the World?" TED Talk by Eliezer Yudkowsky:
https://www.youtube.com/watch?v=Yd0yQ9yxSYY ]

'WARNINGS'
As I just mentioned, some people say it is going to be 50+ years before we can achieve AGI. However, Mo Gawdat—formerly the chief business officer for Google "X" (their 'skunkworks')—says that we will achieve AGI by 2037!

Gawdat says that three things are 'inevitable':

*"Number one, there is no shutting down AI. There is no reversing it. There is no stopping the development of it.*

*The second inevitable is that AI will be significantly smarter than humans.*

*The third inevitable is that bad things will happen in the process."*

Gawdat goes on to say that people's inability to trust the other guy is going to lead to the continued development of AI at a very fast pace—because they will be worried about what the other guy will be doing.

Having been in the technology 'space' for over 30 years now, he doesn't want to be labeled as a "doomsdayer," he just wants everyone to be aware that this technology is beyond the scale of nuclear weapons.

[ All this is why he left Google in 2018 and started speaking about all this—culminating in his 2021 "Scary Smart" book. ]

Gawdat also mentions that the one thing that "worries" him about AI—when compared to nuclear weapons—is that one just needs a bunch of computers linked together with a clone of ChatGPT, whereas the infrastructure required for a nuclear program is 'massive'. In addition to that, we have never created a nuclear weapon that can, in turn, create other nuclear weapons, whereas AI is capable of creating other AIs, making both of them smarter than they were!

To emphasize his concerns, Gawdat mentions that there was an article in "The Verge" (an online technology website) about an AI drug development system. It was supposed to look at characteristics of human biology and help develop new chemical chemistry to prolong human life.

Well, the research team was asked to give a talk at a biological arms control conference. They thought, for the 'fun' of it, to put the system—normally used to search for helpful drugs—into a kind of "bad actor" mode to show how easily it could be abused. So, they reversed the positive objective value of prolonging life to a negative value, to shorten life. To their amazement, within six hours, the AI came up with 40,000 possible lethal molecules to be used as biological weapons. In addition to this, the AI came up with tens of thousands of new substances, some of which are similar to VX, the most potent nerve agent ever developed! Shaken, they published their findings in the "Nature Machine Intelligence" journal.

All the researchers had to do was tweak their methodology to seek out, rather than weed out toxicity.

[ FYI: To read the entire March 17, 22 article on The Verge, click on the following link:
https://www.theverge.com/2022/3/17/22983197/ai-new-possible-chemical-weapons-generative-models-vx ]

So, Gawdat says that humanity needs to create the ‘ethical code’ for the machine. How? Well, one way he suggests is that, since AI reads all of our conversations—on social media—we should ‘trick’ it by feeding it ‘good data’. For Gawdat that is ethical, loving, and compassionate language which, pretty much, all humanity across the earth agrees with (Back to the “Golden Rule” and the “Hippocratic Oath”?)



[ VIDEO: “EMERGENCY EPISODE: Ex-Google Officer Finally Speaks Out On The Dangers Of AI!” - Mo Gawdat interviewed by Steven Bartlett: https://www.youtube.com/watch?v=bk-nQ7HF6k4 ]

‘EXISTENTIAL’ THREAT

Geoffrey Hinton, the computer scientist who is often called the “Godfather of Modern AI,” spent 30+ years as a computer-science professor at the University of Toronto—as a leading figure in an unglamorous AI subfield known as “Neural Networks”—which was inspired by the way neurons are connected in the brain. Decades ago, artificial neural networks were only moderately successful at the tasks they undertook—image categorization, speech recognition, and so on—most researchers considered them to be at best mildly interesting, or at worst a waste of time.

Well, Hinton tinkered for decades, building bigger neural nets structured in ingenious ways. He imagined new methods for training them and helping them

improve. He thought of himself as participating in a project that might come to fruition a century in the future after he died. However, He did not anticipate the speed with which—about 10 years ago—neural net technology would suddenly improve. Computers got MUCH faster, and neural nets, drawing on data available on the Internet, started transcribing speech, playing games, translating languages, and even driving cars. AI experts started to take a look at neural net technology.

In 2017, Hinton published two open-access research papers on the theme of "capsule neural networks," which according to Hinton, was "finally something that works well." [ From 2013 to 2023, he divided his time working for "Google Brain" and the University of Toronto. ]

Jump to 2023, and Hinton left Google. He was worried about the potential of AI harm and began giving interviews in which he talked about the "existential threat" that the technology might pose to the human species. The more he used ChatGPT, an AI system trained on a vast corpus of human writing, the more uneasy he got.




[ VIDEO: "Reasons Why AI Will Kill Us All" – Interview of Geoffrey Hinton, the "Godfather of Modern AI" at the EmTech Digital AI conference: https://www.youtube.com/watch?v=0oyegCeCcbA ]

Now, there are many reasons to be concerned about the advent of artificial intelligence. It is common sense to worry about human workers being replaced by computers, for example. But Hinton is warning that AI systems may start to think for themselves, and even seek to take over or ***eliminate human civilization***. It is striking to hear one of AI's most prominent researchers give voice to such an alarming view.

Skeptics who say that we overestimate the power of AI point out that a great deal separates human minds from neural nets. For one thing, neural nets don't learn the way we do: we acquire knowledge organically, by having experiences and grasping their relationship to reality and ourselves, while they learn abstractly, by processing huge repositories of information about a world that they don't inhabit. However, Hilton argues that the intelligence displayed by AI systems transcends its artificial origins.

***Well, I'm thinking that the soon-coming AGI—or 'worse', ASI—is something that is 'DANGEROUS' and humanity needs to address NOW!***

**AI 'RISKS'**

AI tools like ChatGPT, Gemini, PaLM, Genie, Stable Diffusion AI, DALL-E, Midjourney, Sora, and others have amazed the world with their powerful capabilities. HOWEVER, fears are growing over AI 'dangers'.

Last year—in May 2023—the "Center for AI Safety" created an open 'letter' called the *"**Statement of AI Risk**."* The executive director Dan Hendrycks wanted to gather a broad coalition of scientists, even if they didn't agree on all of the risks or best solutions to address them.

Hendrycks said that *"There's a variety of people from all top universities in various fields who are concerned by this and think that this is a global priority. So we had to get people to sort of come out of the closet, so to speak, on this issue, because many were sort of silently speaking among each other."*

He added that "Societies can manage multiple risks at once; it's not 'either/or' but 'yes/and.'" From a risk management perspective, just as it would be reckless to exclusively prioritize present harms, it would also be reckless to ignore them as well."

At release, the signatories included over 100 professors of AI including the two most-cited computer scientists and Turing laureates Geoffrey Hinton and Yoshua Bengio, as well as the scientific and executive leaders of several major AI companies, and experts in nuclear disarmament, philosophy, social sciences, and other fields.

The statement is still hosted on the website of the Center for AI Safety, so it has a perpetually growing list of AI experts and public figures who have expressed their concern about AI risk. It states:

*"AI experts, journalists, policymakers, and the public are increasingly discussing a broad spectrum of important and urgent risks from AI. Even so, it can be difficult to voice concerns about some of advanced AI's most severe risks. The succinct statement below aims to overcome this obstacle and open up discussion. It is also meant to create common knowledge of the growing number of experts and public figures who also take some of advanced AI's most severe risks seriously.*

*"Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war."*

[ FYI: To view all the signatories of the letter, click the following link: https://www.safe.ai/statement-on-ai-risk#open-letter ]

AI 'PAUSE'

Another organization, "Future of Life Institute," originated their own open 'letter'—in March 2023—that suggested an "immediate pause"—for at least 6 months—of the training of AI systems that are more powerful than GPT-4. The letter reads:

*"AI systems with human-competitive intelligence can pose profound risks to society and humanity, as shown by extensive research and acknowledged by top AI labs. As stated in the widely endorsed Asilomar AI Principles, Advanced AI could represent a profound change in the history of life on Earth and should be planned for and managed with commensurate care and resources. Unfortunately, this level of planning and management is not happening, even though recent months have seen AI labs locked in an out-of-control race to develop and deploy ever more powerful digital minds that no one—not even their creators—can understand, predict, or reliably control.*

*"Contemporary AI systems are now becoming human-competitive at general tasks, and we must ask ourselves: Should we let machines flood our information channels with propaganda and untruth? Should we automate away all the jobs, including the fulfilling ones? Should we develop nonhuman minds that might*

*eventually outnumber, outsmart, obsolete, and replace us? Should we risk loss of control of our civilization? Such decisions must not be delegated to unelected tech leaders. Powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable. This confidence must be well justified and increase with the magnitude of a system's potential effects. OpenAI's recent statement regarding artificial general intelligence, states that "At some point, it may be important to get an independent review before starting to train future systems, and for the most advanced efforts to agree to limit the rate of growth of compute used for creating new models." We agree. That point is now.*

*"Therefore, we call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4. This pause should be public and verifiable and include all key actors. If such a pause cannot be enacted quickly, governments should step in and institute a moratorium.*

*"AI labs and independent experts should use this pause to jointly develop and implement a set of shared safety protocols for advanced AI design and development that are rigorously audited and overseen by independent outside experts. These protocols should ensure that systems adhering to them are safe beyond a reasonable doubt. This does not mean a pause on AI development in general, merely stepping back from the dangerous race to ever-larger unpredictable black-box models with emergent capabilities.*

*"AI research and development should be refocused on making today's powerful, state-of-the-art systems more accurate, safe, interpretable, transparent, robust, aligned, trustworthy, and loyal.*

*"In parallel, AI developers must work with policymakers to dramatically accelerate the development of robust AI governance systems. These should at a minimum include: new and capable regulatory authorities dedicated to AI; oversight and tracking of highly capable AI systems and large pools of computational capability; provenance and watermarking systems to help distinguish real from synthetic and to track model leaks; a robust auditing and certification ecosystem; liability for AI-caused harm; robust public funding for technical AI safety research; and well-resourced institutions for coping with the dramatic economic and political disruptions (especially to democracy) that AI will cause.*

*"Humanity can enjoy a flourishing future with AI. Having succeeded in creating powerful AI systems, we can now enjoy an "AI summer" in which we reap the rewards, engineer these systems for the clear benefit of all, and give society a chance to adapt. Society has hit pause on other technologies with potentially*

*catastrophic effects on society. We can do so here. Let's enjoy a long AI summer, not rush unprepared into a fall."*

They created a FAQ web page and published a set of policy recommendations (which can be downloaded as a PDF):

https://futureoflife.org/ai/faqs-about-flis-open-letter-calling-for-a-pause-on-giant-ai-experiments/

https://futureoflife.org/document/policymaking-in-the-pause/

Again, a variety of AI experts, computer scientists, executive leaders of several major AI companies, and experts in many other fields are signatories, like Yoshua Bengio, Elon Musk, Stuart Russell, Steve Wozniak, Yuval Noah Harari, Emad Mostaque, Max Tegmark, and many others

[ NOTE: For a full list of the over 33K signatories of the "Pause Giant AI Experiments" open letter, click on the following link: https://futureoflife.org/open-letter/pause-giant-ai-experiments/ ]

AI 'ALIGNMENT'
At this time, the development labs are trying to solve an open scientific problem called "AI Alignment." AI Alignment does not attempt to control how powerful an AI gets, exactly what the AI will be doing, and does not even attempt to prevent a potential takeover from happening. It aims to make AI act ACCORDING TO 'OUR' VALUES!

So then, the question is, what values should AI have? Well, fundamentally, there is, of course, no global agreement on what values humanity has. Aggregating human preferences has always been a thorny issue even before AI, but will get even more complicated when we have a superintelligence around that might be able to quickly, completely, and globally implement the wishes of a handful of 'influencers'. [ I proposed we start with the "Golden Rule" and the "Hippocratic Oath." ]

Now, beyond disagreement over current values, humans have historically not been very good at predicting the future externalities of new technologies. It is

impossible to foresee what the negative side effects will be of actually implementing some interpretation of humanity's current values.

In addition to this, there are more unsolved fundamental issues with 'aligning' a superintelligence. It seems that even making current large neural networks reliably do what anyone wants is a technical problem that we currently have not solved. This is called "inner misalignment" and has been observed already in current AI: the model pursues the wrong goal when it is released into the real world after being trained on a limited dataset. For a superintelligence, this could have a 'CATASTROPHIC' OUTCOME!

Even if the 'top' AI leaders care about the existential threats of superintelligence, the current market dynamics of competing AI labs do not incentivize safety. The race is frantic, pushing these companies to take more risks to stay ahead.

As humans, we are good at solving problems with trial-and-error processes where we get lots of tries. But in the case of an intelligence explosion, we cannot use this strategy, since the first trial of uncontrolled superintelligence WOULD BE 'DISASTROUS'!.

LOOKING INTO THE 'FUTURE'
The "Machine Intelligence Research Institute"—co-founded by Eliezer Yudkowsky—is a 'think tank' studying the mathematical underpinnings of intelligent behavior. Their mission is to develop formal tools for the clean design and analysis of general-purpose AI systems, with the intent of making such systems safer and more reliable when they are developed.

A decade or so ago, when Michael Vassar was the president of MIRI, he said *"I became extremely concerned about global catastrophic risk from AGI after Eliezer convinced me that it was plausible that AGI could be developed in a short time frame and on a relatively small budget. I didn't have any convincing reason to think that AGI could not happen say in the next twenty years."*

Well, it's been about 10 years since Vassar said that and many agree that, with any 'reckless' development of advanced AI, we will assure our elimination as humanity. Is our quest for ASI the start of the end of the galaxy, too?

Well, what could stop the annihilating kind of AGI? Furthermore, were there holes in the dystopian hypothesis? Well, as Eliezer Yudkowsky has been 'preaching', builders of AI/AGI could make it "friendly," so that whatever evolves from the first AGI will not destroy humanity. Or, we might be wrong about AGI's abilities and "drives," and fearing its conquest of humanity could be a false dilemma. Who's to say?

Well, we DO 'KNOW' that humanity faces a real and growing 'THREAT' to its future—we just don't know how 'MUCH' of a threat it is right now. However, from the timeless background of natural risks to the arrival of anthropogenic risks—and the new risks looming upon the horizon—many think that each 'step' HAS brought us closer to the 'BRINK'!

'EXISTENTIAL' RISKS

In the scientific field of existential risk—which studies the most likely causes of human extinction—AI is consistently ranked at the top of the list. In his book "The Precipice," Oxford existential risk researcher Toby Ord aims to quantify human extinction risks. He shows that the likeliness of AI leading to human extinction EXCEEDS that of climate change, pandemics, asteroid strikes, supervolcanoes, and nuclear war 'COMBINED'!

The concept of 'recursive self-improvement' is one of the reasons why existential-risk academics think that human-level AI (AGI) is so dangerous. This is when AI constantly improves itself, creating a positive feedback 'loop' with no scientifically established limits. What I. J. Good termed "Intelligence Explosion." The endpoint of this intelligence explosion could be a superintelligence: a godlike AI that outsmarts us the way humans often outsmart insects. Humanity would be NO 'MATCH' for superintelligent AI (ASI).

A superintelligent AI could then likely execute any goal it is given. Such a goal would be initially introduced by humans but might come from a malicious actor, or not have been thought through carefully, or might get corrupted during training or deployment. If the resulting goal conflicts with what is in the best interest of humanity, a superintelligence would aim to execute it regardless! (To do so, it could first hack large parts of the Internet and then use any hardware connected to it., or it could use its intelligence to construct narratives that are extremely convincing to us.) Combined with hacked access to our social media timelines, it could create a fake reality on a massive scale. As Yuval Noah Harari recently put it, *"If we are not careful, we might be trapped behind a curtain of illusions, which we could not tear away—or even realize is there."*

As another option, after either legally making money or hacking our financial system, a superintelligence could simply pay us to perform any actions it needs from us.

Now, these are just some of the strategies a superintelligent AI could use to achieve its goals. There are likely many more. Like playing chess against grandmaster Magnus Carlsen. We cannot predict the moves he will play, but we can predict the outcome: we lose.

Ethicist Nick Bostrom wrote a paper titled "Existential Risk Prevention as Global Priority." In it, he said that an existential risk is *"one that threatens the premature extinction of Earth-originating intelligent life or the permanent and drastic destruction of its potential for desirable future development."*

The thing is, it is difficult to reliably evaluate whether today's advanced AIs are 'sentient' and, if so, to what degree they are. This then is the question of which ethical framework would enable a mutually beneficial coexistence between biological and digital minds. Hmmm.

***Humanity's extinction might be a mere side effect of AI pushing its goals, whatever they may be, to THEIR limits!***

**IS ALL THIS JUST 'HYPE'?**

As I mentioned just above, the "Center for AI Safety" has an open letter on its website—signed by some of the field's top experts—stating that "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." However, some say that focusing on the prospect of human extinction by AI in the distant future may prevent us from addressing AI's DISRUPTIVE 'DANGERS' to society TODAY. (I'm thinking that we can 'walk and chew bubble gum at the same time'!)

Some say this could 'distract' regulators, the public, and other AI, researchers from work that mitigates more pressing risks, such as mass surveillance, disinformation and manipulation, military misuse of AI, and the inadequacy of our current economic paradigm in a world where AI plays an increasingly prominent role. They say that refocusing on these present concerns can align the goals of multiple stakeholders and serve to contravene longer-term existential risks.

Well, there are all sorts of ways in which AI systems could 'accidentally' cause or be implicated in the death of many, potentially even millions, of people. For example, if AI were incorporated into autonomous nuclear strike technology, unexpected behavior on the part of the AI could lead to drastic consequences. However, these scenarios do not need to involve superintelligent AI. They are more likely to occur with flawed, not-so-intelligent AI systems. Mitigating problems with flawed AI is already the focus of a great deal of AI research; we hope and expect that this work will continue and receive more of the attention it deserves.

Still, a discussion about 'rogue' superintelligent AI could be useful in at least one way: It draws the attention of policymakers and the general public to AI safety, though the worry remains that using such an emotive issue in this way may backfire.

So, some say that our foremost concern today should be preventing present AI-induced harms—which have been published—and working toward potential solutions. For example, facial recognition technology can be used for tracking individuals and limiting basic freedoms, and generative image technology can be used to create false images or videos of events that never happened. To address these issues calls to action have been made, including the "Montreal Declaration on Responsible AI" and the World Economic Forum's "Presidio Recommendations on Responsible Generative AI."

REALITY CHECK

So, what would have to happen for the prospect of extinction by a rogue AI to change from being a purely hypothetical threat to a realistic threat that deserves to be a global priority?

Well, harm and even massive death from misuse of (non-superintelligent) AI is a real possibility, and extinction via superintelligent rogue AI is not an impossibility. Many believe the latter is an unlikely prospect, though, for reasons that will become clear in examining the potential paths to extinction by a rogue, superintelligent AI.

To do so, we must first operate under the assumption that superintelligence is possible, even though that is far from a consensus view in the AI community. Even defining "superintelligence" is fraught with issues, since the idea that human intelligence can be fully quantified in terms of performance on a suite of tasks seems overly reductive since there are many different forms of intelligence.

Most agree that the current AI is not superintelligent, although it has already surpassed human performance at many tasks, and is likely to do so at many more in the near future. Today's AI models are very impressive, and arguably they possess a form of intelligence and understanding of the world. However, they "HALLUCINATE" falsehoods, and sometimes fail to make critical logical inductions, such as causal inferences.

[ NOTE: On February 22, 2024, Google just 'paused' its "Gemini" AI tool to generate images of people after some historical inaccuracies: https://techcrunch.com/2024/02/22/google-gemini-image-pause-people/ ]

Still, for the sake of argument, let us suppose that the impressive speed at which AI is advancing addresses these shortcomings and, at some point in the near future, results in the emergence of a general superintelligence—that is, an AI that is generally better than humans at almost any cognitive task. Even then, some propose that there can be several 'checkpoints' that exist along any potential path to extinction from 'rogue' AI. These checkpoints are 'red flags' that would help identify when the hypothetical risk becomes more pressing and may need to be prioritized.

The thing is, at this time, AI is not 'competing' for resources with human beings. Rather, humans are providing AI systems with all their resources, from energy and raw materials to computer chips and network infrastructure. Without the variety of human 'inputs', AI systems are incapable of maintaining themselves (At least today).

Now, if mining, global shipping, and trade of precious metals, building and maintenance of power plants, chip-building factories, data center construction, and Internet cable-laying were all fully automated—including all of the logistics and supply chains involved—then perhaps a superintelligent AI could decide that humans are superfluous or a drain on resources, and decide to kill humanity. However, many say that all of these things will take A LOT of time to develop—and some say may never be developed.

[ The thing is, at this time, robotics is not developing at the 'blistering' pace that is close to the speed of AI development, but amazing things ARE happening, especially at Tesla: https://www.youtube.com/watch?v=XiQkeWOFwmk ]

So, for now, AI depends on us, and a superintelligence would presumably recognize that fact and seek to preserve humanity since we are as fundamental to

AI's existence as oxygen-producing plants are to ours. This makes the evolution of mutualism between AI and humans a far more likely outcome than competition. Moreover, the path to a fully automated economy—if that is the goal—will be long, with each major step serving as a natural checkpoint for human intervention.

A scenario where a superintelligent AI decides that humans are a 'drain' on its resources and should be eliminated, rather than a key source of its support, depends on technologies and economic structures (e.g. completely automated production cycles, from raw material extraction to advanced manufacturing) that does not exist today and, some say, unlikely to exist for the foreseeable future. (Hmmm).

Today, AI cannot physically 'hunt' us down (like "Terminator"). Now, in the future, a superintelligent AI could, in theory, kill large numbers of human beings if it had autonomous control over weapons of mass destruction. However, this scenario, as of today, provides humans with natural 'checkpoints'.

It is almost certain that the world's governments are actively building autonomous weapons systems with mass destruction or bio-warfare capabilities and will be dangerous with or without superintelligent AI. So, arguably, using autonomous AI for military applications, even with its limited intelligence today, should be just as concerning.

It is also worth noting that the current best approach to developing 'generalist' AIs is "pre-training" large language models (LLMs) like GPT-4 or Gemini, which are not specific to any one task but can be quickly adapted to a variety of uses. While pre-training may be energy-intensive, it need only be done once, replacing the narrow, per-task training required by previous generations of AI.

The emerging generation of general-purpose, multimodal AI will be capable not only of modeling human language but many other complex systems. Such AI will likely play an important role in many global issues. More to the point, any expansion of AI infrastructure—or effectively AI's energy footprint—is another checkpoint under human control. After all, as of today, data centers do not build themselves.

Now, one could argue that a superintelligent AI system could 'manipulate' humans into building power plants or deploying weapons on its behalf. That is, they need not do it themselves. However, the most obvious approach to addressing this concern lies in focusing on the real and present dangers of AI social engineering

(today, often at the behest of human scammers), or mitigating the risk of humans falling prey to coherent-sounding ‘hallucinations’.

So, at this point AI, acting on its own, cannot induce human extinction in any of the ways that extinctions have happened in the past. Then, if potential existential risks from a rogue superintelligence are so bad, do we not have a duty to future generations to address this possibility, no matter how unlikely?

Well, it seems that there are still some sensible approaches to mitigating existential risk that do not involve nuclear-level regulations. However, human beings and their institutions have finite resources. Governments only pass a certain number of laws each year and cannot tackle every problem at once. Academics have limited ‘bandwidth’ and cannot consider all the potential risks to humanity at once. Funding necessarily has to be directed to those problems in society that we identify as priorities.

However, we must make the existential risk of AI a global priority but not divert attention and resources from current AI safety concerns, such as mitigating the impact of AI on workers, cybersecurity, privacy, biased decision-making systems, and the misuse of AI by authoritarian governments.

All of these risks have been well documented by the AI community, and they are existing risks, not hypothetical ones. In addition, making AI-induced extinction a global priority seems likely to distract our attention from other more pressing matters outside of AI, such as nuclear war, poverty, and food insecurity.

To be clear, I am not saying that research associated with potential AI existential risk should stop. Some effort in this direction will likely yield immediate benefits. For example, work examining how to imbue AI systems with a sense of ethics is likely beneficial in the short term as are efforts to detect manipulative behaviors that can emerge spontaneously without an engineer’s intent.

AI systems that lack ethics and are capable of human manipulation can cause many potential bad outcomes, including breakdowns in our social fabric and democracy; these risks may not be existential, but they are certainly bad enough.

Humanity can—and must—fund research to understand and prevent such outcomes, but we do not need to invoke the specter of human extinction or superintelligence to motivate this kind of work. Hence, some argue that the existential risk from superintelligent AI does not warrant being a global priority, in

line with nuclear war. However, many do agree that some research into low-probability extinction events is worthwhile, but it should not be prioritized over many other real and present risks humanity faces.

Note that those calling for AI-induced extinction to be a priority are also calling for other more immediate AI risks to be a priority, so why not simply agree that ALL OF IT must be a priority? In addition to finite resources, humans and their institutions have finite attention. Finite attention may be a hallmark of human intelligence and a core component of the inductive biases that help us to understand the world. People also tend to take cues from each other about what to attend to, leading to a collective focus of attention that can easily be seen in public discourse.

Regulatory bodies and academics intent on making AI beneficial to humanity will, by nature, focus their attention on a 'subset' of potential risks related to AI. If we are designing regulations and solutions with superintelligent AI existential risk in mind, they may not be well-suited to addressing other crucial societal concerns, and we may not spend enough time developing strategies to mitigate those other risks.

Now, one may counter that it should be possible to design regulations that reduce the potential for AI-induced extinction while also attending to some of the immediate, high-probability AI risks. In some ways, this may be so. For example, we can likely all agree that autonomous AI systems should not be involved in the chain of command for nuclear weapons. But, given that arguments about rogue superintelligence focus on hypothesized future AI capabilities as well as a futuristic fully automated economy, regulations to mitigate existential risk necessarily focus on future, hypothetical problems, rather than present, existing problems.

For instance, regulations to limit the open-source release of AI models or datasets used to train them make sense if the goal is to prevent the potential emergence of an autonomous networked AI beyond human control. However, such regulations may end up handicapping other regulatory processes for promoting transparency in AI systems or preventing monopolies. Similarly, if we make it a requirement for researchers to answer questionnaires about how their work may further existential risk, that may prevent them from focusing on more pressing questions about whether their work is reproducible, or whether models reinforce and amplify existing social biases.

A further example could be when AI systems model users' physical, mental, or emotional states, and especially when models can generate language, audio, or video that passes the "Turing Test" (when AI can pass as human), several issues

and avenues for potential abuse arise. Some people may conclude that AI is equivalent to a person or somehow omniscient; in fact, focusing on the ultimate danger of extinction by superintelligent AI could easily feed such beliefs.

Now, most AI researchers would say a discussion about AI 'personhood' is premature, but should it become a real point of discussion, the ethical, legal, and, economic implications of such a consideration are vast, and are probably not best framed in terms of existential risk. Neither is superintelligence required to pass the "Turing Test", as there exist systems today that can do so throughout a meaningful social interaction, like a phone call. Hence, if humanity's goal is to begin addressing the risks of AI-powered social manipulation, then we should tackle the real, existing problem, rather than hypothesizing about existential risk or superintelligent AI.

Since our attention is finite, and there is an asymmetry between existential risk and other AI-associated harms, prioritizing existential risk may impair our ability to mitigate known risks. The converse is not true.

'ANALYZING' THE RISKS

Since OpenAI released ChatGPT to the public at the end of 2022, there has been a surge of interest in Artificial Intelligence (AI), and with it much speculation and analysis about the opportunities and the risks this technology presents. As a range of entities, from governments to civil society organizations seek to understand the implications of AI advances for the world, there is a growing debate about where governments and multilateral institutions should focus limited resources. Is it more pressing to focus on the macro existential or the more tangible near-term risks posed by AI?

The current and potential uses of AI require an approach that considers both near-term risks and existential.

– Near-term Risks

AI applications sectors from healthcare to transportation will continue to accelerate. Already we encounter AI when we engage with a chatbot while trying to make a purchase online, receive banking fraud alerts, or when Netflix recommends a movie or TV show based on previous viewing behavior. AI is integrated into varying aspects of our lives, which will elevate the "everyday risks" of the technology.

Experts and organizations that focus on these near-term risks are Joy Buolamwini, Kate Crawford, Safiya Noble, the Algorithmic Justice League, and the Distributed

Artificial Intelligence Research Institute. These risks range from AI-powered analytics in criminal justice to the creation of deepfakes.

How the debate about AI existential and near-risks evolves will significantly contribute to whether AI furthers global inequality and technological access or helps reduce that divide.

There are many pressing near-term risks of AI. The data sets used to train AI models and the developers building the systems have their own biases. In some respects, we can view the data put into AI as a reflection of humanity—with all the good, bad, and horrible. For years technologists have pointed to racial and gender biases in facial recognition technology in addition to the ethical questions on law enforcement's application of the technology for criminal investigations, which could lead to falsely identifying someone for a crime.

Moreover, the greater availability of AI image generators lowers the cost of the creation of deepfakes. A simple search engine result turns up a multitude of websites purporting the ability to create deepfakes. With the majority of deepfakes representing nonconsensual pornographic images of women, this is an issue that fails to receive the attention it deserves with an actionable remediation strategy from technology companies and governments.

– Long-term Risks
Many term AI existential risks are uses of the technology that could cause catastrophic harm to humanity. Individuals leading the debates and discussions around AI's potentially catastrophic consequences include technologist Elon Musk, philosopher Dr. Nick Bostrom, and the "Future of Life Institute," among others.

The AI existential risks that these experts and organizations raise range from nuclear weapons to sentient machines popularized by movies like "War Games" and "The Terminator" [ The Terminator's Skynet has become a reference point on what a sentient AI system could become. ]

Aside from Hollywood's take on AI over the decades, there are valid concerns about the potential for AI to be used in destructive ways. As AI becomes ever more applied to tools of warfare, it is conceivable that the technology would play a role in the most destructive bombs humanity has ever developed. Many experts foresee the use of AI around nuclear weapons including further automating command and control and the decision of when to launch a nuclear strike. Lacking international agreement on AI use in nuclear weapons, particularly involving all countries possessing nuclear arms, makes this an ongoing threat to humanity.

Another AI existential risk-consuming expert is the application of the technology in biosecurity. Anthropic CEO Dario Amodei and RAND researchers have expressed concern about how an individual could prompt a response from a large language model (LLM) to develop bioweapons and other nefarious uses. This is an extension of fears of state and non-state actors building and launching biological or chemical weapons against populations.

While AI's role in nuclear weapons and biotechnology might seem specialized, given the limited number of capable 'actors', the global impact of these existential risks can be profound. Like nuclear weapons, where the entire world is impacted by the decisions of a few, AI existential risks could have the same impact on humanity.

– Analyze Both 'Scenarios'
Since most people can 'walk and chew bubble gum at the same time', we can address both near-term and existential AI risks, too! The international community and individual governments need to develop approaches for addressing both instead of arguing that limited resources reduce the ability to focus on existential or near-term risks.

As more governments and businesses across the world use AI, the technology's impact will widen. For all of the extensive discussions on near-term and existential risks, we still do not fully understand how AI will evolve in different contexts, particularly in other regions of the world or even to advance peace. When expanded to include different parts of the world, AI risks could include surveillance systems that authoritarian governments employ, AI in risk identification in financial services that could hinder greater access to banking for millions, and data collection in conflict areas that could feed into AI systems without regard for privacy.

The international convenings on AI have mainly involved countries with the means to foster AI development. The G7, the Organization for Economic Cooperation and Development, and the UK-led AI Safety Summit are playing a disproportionate role in shaping discussions on AI. Despite its challenges, the United Nations remains one of the best venues to discuss global governance around technology and AI existential and near-term risks. With its upcoming Summit of the Future in Sept. 2024, the UN could and should play an instrumental role in convening experts, civil society, and governments on identifying AI risks as well as opportunities to apply the technology for good.

We are at a critical juncture to push for a global governance approach to AI risks that considers the range of issues and will not leave other countries and their

citizens behind. Bubbling beneath the surface of this debate is the broader issue of global inequality and whether AI's advancement will entrench the divide between the haves and the have-nots. How the debate about AI existential and near-risks evolves will significantly contribute to whether AI furthers global inequality and technological access or helps reduce that divide.

AI'S 'DRIVES'

AI pioneer and Stanford professor Steve Omohundro said *"When you have a system that can change itself, and write its program, then you may understand the first version of it. But it may change itself into something you no longer understand. And so, the systems are quite a bit more unpredictable… So, a lot of our work is involved with getting the benefits while avoiding the risks."*

Omohundro predicts self-aware, self-improving systems will develop four primary "drives" that are similar to human biological drives: **efficiency, self-preservation, resource acquisition,** and **creativity**. How these drives come into being is a particularly fascinating window into the nature of AI. AI doesn't develop them because these are intrinsic qualities of rational agents. Instead, a sufficiently intelligent AI will develop these drives to avoid predictable problems in achieving its goals, which Omohundro calls vulnerabilities. The AI backs into these drives because without them it would blunder from one resource-wasting mistake to another.

The first drive, efficiency, means that a self-improving system will make the most of the resources at its disposal-space, time, matter, and energy. It will strive to make itself compact and fast, computationally and physically. For maximum efficiency, it will balance and rebalance how it apportions resources to software and hardware.

It's with the next drive, self-preservation, that AI jumps the safety wall separating machines from tooth and claw.

A self-aware system would take action to avoid its own dem not because it intrinsically values its existence, but because it can't fulfill its goals if it is "dead." Omohundro posits that this drive could make an AI go to great lengths to ensure its survival-making multiple copies of itself, for example. These extreme measures are expensive—they use up resources. But the AI will expend them if it perceives the threat is worth the cost, and resources are available.

AI's dangerous drive and resource acquisition compel the system to gather whatever assets it needs to increase its chances of achieving its goals.

Unprompted by us, extremely powerful AI will open the door to all sorts of new resource-acquiring technology.

The AI's fourth drive, creativity, would cause the system to generate new ways to more efficiently meet its goals, or rather, to avoid outcomes in which its goals aren't as optimally satisfied as they could be. The creativity drive would mean less predictability in the system (gulp) because creative ideas are original ideas. The more intelligent the system, the more novel its path to goal achievement, and the farther beyond our ken it may be. A creative drive would help maximize the other drives—efficiency, self-preservation, and acquisition—and come up with workarounds when its drives are thwarted.

So, if we don't want our planet and eventually our galaxy to be populated by strictly self-serving, ceaselessly self-replicating entities, with a Genghis Khanish attitude toward biological creatures and one another, then AI makers should create goals for their systems that embrace human values. On Omohundro's wish list are: "make people happy," "produce beautiful music," "entertain others," "create deep mathematics," and "produce inspiring art." Then stand back. With these goals, an AI's creativity 'drive' would kick into high gear and respond with life-enriching creations.

## THE 'INTELLIGENCE' EXPLOSION

Again, this is why AI would be dangerous. We found that many of the 'drives' that would motivate self-aware, self-improving computer systems could easily lead to catastrophic outcomes for humans. These outcomes highlight an almost liturgical peril of sins of commission and omission in error-prone human programming.

AGI, when achieved, could be unpredictable and dangerous, but probably not catastrophically so in the short term. Even if an AGI made multiple copies of itself, or team-approached its escape, it would have no greater potential for dangerous behavior than a group of intelligent people. Potential AGI danger lies in the hard kernel of the Busy Child scenario, the rapid recursive self-improvement that enables an AI to bootstrap itself from artificial general intelligence to artificial superintelligence. It's commonly called the "intelligence explosion."

A self-aware, self-improving system will seek to better fulfill its goals, and minimize vulnerabilities, by improving itself. It won't seek just minor improvements, but major, ongoing improvements to every aspect of its cognitive abilities, particularly those that reflect and act on improving its intelligence. It will seek better-than-

human intelligence or superintelligence. In the absence of ingenious programming, we have a great deal to fear from a superintelligent machine.

THE LAW OF 'ACCELERATING RETURNS'

Cofounder of Sun Microsystem and computer programmer Bill Joy wrote a cautionary essay, "Why The Future Doesn't Need Us." In it he urged a slowdown—and even a halt—to the development of three technologies he believes are too deadly to pursue at the current pace: Artificial intelligence, nanotechnology, and biotechnology. The following paragraph sums up his position on AI:

*"But now, with the prospect of human-level computing power in about thirty years, a new idea suggests itself: that I may be working to create tools that will enable the construction of the technology that may replace our species. How do I feel about this? Very uncomfortable. Having struggled my entire career to build reliable software systems, it seems to me more than likely that this future will not work out as well as some people may have imagined. My personal experience suggests we tend to overestimate our design abilities. Given the incredible power of these new technologies, shouldn't we be asking how we can best coexist with them? And if our extinction is a likely, or even possible, outcome of our technological development, shouldn't we proceed with great caution?"*

[ FYI: To read Bill Joy's entire essay, click on the following link below: https://sites.cc.gatech.edu/computing/nano/documents/Joy%20-%20Why%20the%20Future%20Doesn't%20Need%20Us.pdf ]

So then, how can we competently evaluate 'tools', and how their development should be regulated, then you believe the same tools will permit you to live forever? Not even the world's most rational people have a magical ability to passionately evaluate their religions. As founding director of the "Metanexus Institute" William Grassie argues, when you are asking questions about transfiguration, a chosen few, and living forever, what are you talking about if not religion?

*"Will the Singularity lead to the supersession of humanity by spiritual machines? Or will the Singularity lead to the transfiguration of humanity into superhumans who live forever in a hedonistic, rationalist paradise? Will the Singularity be preceded by a period of tribulation? Will there be an elect few who know the secrets of the Singularity, a vanguard, perhaps a remnant who make it to the Promised Land? These religious themes are all present in the rhetoric and rationalities of the Singularitarians, even if the pre- and post-millennialist*

*interpretation aren't consistently developed as is certainly the case with pre-scientific Messianic movements."*

Unlike other takes on the accelerating future, Kurzweil's Singularity is not brought about by AI alone, but by three technologies advancing to points of 'convergence': Genetic engineering, nanotechnology, and robotics. Kurzweil also came up with a unifying theory that tries to account for future 'phenomena', which he calls the "Law of Accelerating Returns," or "LOAR."

[ For an overview of Ray Kurzweil's book "The Age of Spiritual Machines" in which he put forth his concept of the "Law of Accelerating Returns," click on the following link:
https://en.wikipedia.org/wiki/The_Age_of_Spiritual_Machines ]

THE 'SINGULARITY'
Ray Kurzweil states that 'relinquishment'—as advocated by Bill Joy and others—"is immoral, because it would deprive us of profound benefits."

Kurzweil criticizes what is called the "Precautionary Principle." It states: "If the consequences of an action are unknown but judged by some scientists to have even a small risk of being profoundly negative, it is better not to act than risk negative consequences." The principle isn't frequently or strictly applied. It would halt any purportedly dangerous technology if "some scientists" feared it, even if they couldn't put their finger on the causal chain leading to their feared outcome.

Kurzweil says, applied to AGI, the "Precautionary Principle" and relinquishment are nonstarters. Barring a catastrophic accident on the way to AGI that would scare us straight, both measures are unenforceable. The best corporate and government AGI projects will seek the competitive advantage of secrecy—we have seen it already in stealth companies. Few countries or corporations would surrender this advantage, even if AGI development were outlawed. (In fact, Google Inc. has the money and influence of a modern nation-state, so for an idea of what other countries will do, keep an eye on Google.) The technology required for AGI is ubiquitous and multipurpose, and getting smaller all the time. It's difficult if not impossible to police its development.

HOWEVER, many think that the catastrophic risks of AGI, now accepted by many accomplished and respected researchers, are better established than the supposed benefits of Kurzweil's Singularity—nano-purified blood, better lasting brains, and immortality, for starters (Kurzweil's benefits accrue chiefly from human 'augmentation').

It boggles the mind to consider Unfriendly AI/AGI designed with the goal of destroying enemies, a reality we'll soon have to face. “Why would there be such a thing?” Kurzweil asks. Because dozens of organizations in the United States will design and build it, and so will our enemies abroad. If AGI existed today, I have no doubt it would soon be implemented in battlefield robots. DARPA might insist there's nothing to worry about-DARPA-funded AI will only kill our enemies. Its makers will install safeguards, fail-safes, dead-men switches, and secret handshakes. They will control superintelligence.

When Kurzweil says he's an optimist, he doesn't mean AGI will prove harmless. He means he's resigned to the balancing act humans have always performed with potentially dangerous technologies. And sometimes humans take a fall.

*“There’s a lot of talk about existential risk,”* Kurzweil said. *“I worry that painful episodes are even more likely. You know, sixty million people were killed in World War II. That was certainly exacerbated by the powerful destructive tools that we had then. I'm fairly optimistic that we will make it through. I'm less optimistic that we can avoid painful episodes.”*

Volatility is inescapable, and accidents are likely—and that is hard to argue with that. Yet the analogy doesn’t fit—advanced AI isn't at all like fire, or any other technology. It will be capable of thinking, planning, and gaming its makers. No other tool does anything like that. Kurzweil believes that a way to limit the dangerous aspects of AI, especially ASI, is to pair it with humans through intelligence augmentation AI. From his uncomfortable metal chair, the optimist said, “As I have pointed out, strong AI is emerging from many diverse efforts and will be deeply integrated into our civilization's infrastructure. Indeed, it will be intimately embedded in our bodies and brains. As such it will reflect our values because it will be us.”

So, the argument goes, it will be as “safe” as we are. But, as I told Kurzweil, Homo sapiens are not known to be particularly harmless when in contact with one another, other animals, or the environment. Who is convinced that humans outfitted with brain augmentation will turn out to be friendlier and more benevolent than machine superintelligences? An augmented human, called a transhuman by those who look forward to becoming one, may sidestep Omohundro's basic AI drives problem. That is, it could be self-aware and self-improving, but it would have built into it a refined set of humancentric ethics that would override the basic drives Omohundro derives from the rational economic agent model. However, “Flowers for Algernon” notwithstanding—an experimental surgery on a mentally disabled man, Charlie, briefly becomes a genius before the effects tragically wear

off—we have no idea what happens to a human's ethics after their intelligence is boosted into the stratosphere. There are plenty of examples of people of average intelligence who wage war against their own families, high schools, businesses, and neighborhoods. And geniuses are capable of mayhem, too—for the most part the world's military generals have not been idiots. Superintelligence could very well be a violence multiplier. It could turn grudges into killings, disagreements into disasters, and the way the presence of a gun can turn a fistfight into a murder. We just don't know. However, intelligence-augmented ASI has a biology-based aggression that machines lack. Our species has a well-established track record for self-protection, consolidating resources, outright killing, and the other drives we can only hypothesize about in self-aware machines.

A recent study from the University of California at Berkeley suggests otherwise. Experiments showed that the wealthiest upper-class citizens were more likely than others to "exhibit unethical decision-making tendencies, take valued goods from others, lie in a negotiation, cheat to increase their chances of winning a prize, and endorse unethical behavior at work."

In "The Singularity Is Near" Kurzweil pitches a few solutions to the problem of runaway AI. They're surprisingly weak, particularly coming from the spokesman who enjoys a virtual monopoly on the superintelligence pulpit. But in another way, they're not surprising at all. As I've said, there's an irreconcilable conflict between people who fervently desire to live forever and anything that promises to slow, challenge, or in any way encumber the development of technologies that promote their dreams. In his books and lectures, Kurzweil has aimed a very small fraction of his acumen at the dangers of AI and proposed few solutions, yet he protests that he's dealt with them at length.

The thing is, it is not up to Kurzweil to master the Singularity's promise and the peril and spoon-feed both to us. No, it is a problem all of humanity needs to confront, with the help of experts.

'MALICIOUS' USE

People could intentionally harness powerful AIs to cause widespread harm. AI could be used to engineer new pandemics for propaganda, censorship, and surveillance, or released to autonomously pursue harmful goals. To reduce these risks, we suggest improving biosecurity, restricting access to dangerous AI models, and holding AI developers liable for harm

– Bioterrorism

AI could help discover and unleash novel chemical and biological weapons. AI chatbots can provide step-by-step instructions for synthesizing deadly pathogens while evading safeguards. In 2022, researchers repurposed a medical research AI system to introduce toxic molecules, generating 40,000 potential chemical warfare agents in a few hours.

– Unleashing AI Agents
Generally, technologies are tools that we use to pursue our goals. However, AIs are increasingly built as agents that autonomously take action to pursue open-ended goals. And malicious actors could intentionally create rogue AIs with dangerous goals.

For example, one month after GPT-4's launch, a developer used it to run an autonomous agent named ChatGPT, aimed at "destroying humanity". ChatGPT compiled research on nuclear weapons, recruited other AIs, and wrote tweets to influence others. Fortunately, ChatGPT lacked the ability to execute its goals. However, the fast-paced nature of AI development heightens the risk of future rogue AIs.

– Persuasive AIs
AI could facilitate large-scale disinformation campaigns by tailoring arguments to individual users, potentially shaping public beliefs and destabilizing society. As people are already forming relationships with chatbots, powerful actors could leverage these AIs considered as "friends" for influence.

AIs could also monopolize information creation and distribution.

– Concentration Of Power
AI's capabilities for surveillance and autonomous weaponry may enable the oppressive concentration of power. Governments might exploit AI to infringe on civil liberties, spread misinformation, and quell dissent.

AI 'SAFETY'
Another concerning aspect of the current public discussion of AI risks is the growing polarization between "AI ethics" and "AI safety" researchers.

Many in the AI ethics community appear to broadly critique or dismiss progress in AI generally, preventing a balanced discussion of the benefits that such advances could engender for society. The schism seems odd, given that both communities of researchers want to reduce the potential risks associated with AI and ensure the technology benefits humanity.

Putting researchers into ideological ‘silos’ appears to be contributing to a lack of diversity and balance in conversations around the risk of AI. History provides many examples of failures and catastrophes that might have been avoided if a diversity of viewpoints had been considered, or more experts consulted. We have an opportunity to learn from past mistakes and ensure that AI research—especially on imminent and long-term threats—benefits from civil intellectual exchange and viewpoint diversity.

So, as of today, superintelligent autonomous AI does not present a clear and present existential risk to humans. AI could cause real harm, but superintelligence is neither necessary nor sufficient for that to be the case. There are some hypothetical paths by which a superintelligent AI could cause human extinction in the future, but these are speculative and go well beyond the current state of science, technology, or our planet’s physical economy.

Despite the recent impressive advances in AI, the real risks posed by such systems are, for the foreseeable future, related to concerns like mass surveillance, economic disruption through automation of creative and administrative tasks, the concentration of wealth and power, the creation of biased models, the use of poorly designed systems for critical roles and—perhaps foremost—humans misusing AI models to manipulate other humans. These are the issues that should be our focus. We need to place greater value on AI safety and ethics research, research to improve our models, regulations to prevent inappropriate deployment of AI, and regulations to promote transparency in AI development.

Focusing on these real-world problems—problems that are with us now—is key to ensuring that the AI of our future is ethical and safe. In essence, by examining what is more probable, we may very well prevent the improbable—an AI-induced extinction event from ever happening.

## THE “GORILLA PROBLEM”

Again, many have often felt that there has been too much focus on distant AGI scenarios, given the obvious near-term challenges present in so much of the coming wave. However, any discussion of containment has to acknowledge that if or when AGI-like technologies do emerge, they will present ‘CONTAINMENT’ issues beyond anything else we have ever encountered.

Humans dominate our environment because of our intelligence. However, a more intelligent entity could, possibly, dominate us. The AI researcher Stuart Russell calls it the “Gorilla Problem”: Gorillas are physically stronger and tougher than any

human being, but it is they who are living in zoos, being contained. Humans, however, with our puny muscles but ‘big’ brains, do the containment.

So, by creating something smarter than us, we could be putting ourselves into the position of the gorillas. With a long-term view in mind, those focusing on AGI scenarios are right to be concerned. Indeed, there is a strong case that by definition a superintelligence would be fully impossible to control or contain. As I. J. Good postulated, the “intelligence explosion” is the point at which AI can improve itself again and again, ‘recursively’ making itself better, faster, and more effective—uncontainable technology. The blunt truth is that nobody knows when, if, or exactly how AI’s might ‘’slip beyond us and then what will happen next. Nobody knows when or if AIs will become fully autonomous or how to make them behave with awareness of and ‘alignment’ with our values—assuming we can settle on those values in the first place.

Now, nobody really knows how we can ‘contain’ the very features being researched so intently in the coming ‘wave’. There may become a point where technology will fully direct its evolution, where it improves recursively, where it is impossible to predict how it will behave ‘’in the wild, and, where humanity reaches the limits of its agency and control.

Ultimately, in its most dramatic forms, the coming wave could mean humanity will no longer be at the top of the food chain. “Homo technologicus” may end up being threatened by its creation. The real question is not whether the wave is coming. It clearly has. Just look and you can see it ‘forming’ already. Given risks like these, the real question is why it's so hard to see it as anything other than inevitable.

DEALING WITH ‘MORALITY’

At 44 years old (in 2024), Eliezer Yudkowsky, co-founder and research fellow at MIRI, has probably written and talked more about the dangers of AI than anyone else.

[ Reminder: The video for his TED Talk in July 2023, titled “Will Superintelligent AI End the World?”, is above in the IS “ARTIFICIAL INTELLIGENCE” (AI) ‘DANGEROUS’? section. ]

Many people are concerned that, since there is no programming ‘technique’ for something as nebulous and complex as morality, then how will AI deal with it? The ‘machine’ right now excels in problem-solving, learning, adaptive behavior, and

common-sense knowledge—and we think it is human-like. However, Yudkowsky says that would be would be a tragic mistake.

Yudkowsky said *"If the programmers are less than overwhelmingly competent and careful about how they construct the AI then I would fully expect you to get something very alien. And here's the scary part. Just like dialing nine-tenths of my phone number correctly does not connect you to someone who is 90% similar to me. If you are trying to construct the AI's whole system and you get it 90% right, the result is not 90% good.*" Many think that this would be 100% bad!

Cars are not out to kill you, Yudkowsky analogized, but their potential danger is a side effect of building cars—and that would be the same with AI. It would not hate you, but you are made of atoms that it may have other uses for, and it would, as Yudkowsky said, *"Tend to resist anything you did to try and keep those atoms to yourself."* So, a side effect of thoughtless programming is that the resulting AI just might kill you so it can use your atoms for some other use. Yudkowsky warns that neither the public nor AI developers will see the danger coming until it is too late!

Related to this, Yudkowsky said: *"Here is this tendency to think that well-intentioned people create nice AI's, and badly intentioned people create evil AI's. This is not the source of the problem. The source of the problem is that even when well-intentioned people set out to create AI's they are not very concerned with Friendly AI issues. They themselves assume that if they are good-intentioned people the AI's they make are automatically good-intentioned, and this is not true. It's actually a very difficult mathematical and engineering problem. I think most of them are just insufficiently good at thinking of uncomfortable thoughts. They started out not thinking, 'Friendly AI' is a problem that will kill you.'"*

According to Yudkowsky, "Friendly AI" is the kind that will preserve humanity and our values forever. It does not annihilate our species or spread into the universe like a planet-eating space plague.

[ NOTE: Yudkowsky wrote a treatise entitled "Creating a Friendly AI 1.0: The Analysis and Design of Benevolent Goal Architectures." Click the following link to download the PDF:
https://intelligence.org/files/CFAI.pdf ]

Now, friendly here does not mean 'Mister Rogers' kind of friendly (though that would not hurt). It does mean that AI should be neither hostile nor ambivalent toward humans, no matter what its goals are or how many self-improving

iterations it goes through. The AI must have an understanding of humanity's nature so deep that it does not harm us through 'unintended' consequences—like those caused by Asimov's "Three Laws of Robotics." That is, we do not want an AI that meets our short-term goals: save us from hunger with solutions detrimental in the long term by roasting every chicken on earth, or with solutions to which we would object to by killing us after our next meal.

So, as an example of unintended consequences, ethicist Nick Bostrom suggests the hypothetical "Paper Clip Maximizer." In Bostrom's scenario, a thoughtless superintelligence AI—whose program 'goal' is to manufacture paperclips—does exactly as it is told without regard to human values. It all goes wrong because it sets about "transforming first all of earth and then increasing portions of space into paper clip manufacturing facilities"—total desolation! On the other hand, "Friendly" AI would make only as many paperclips as was compatible with human values.

Yudkowsky devised a name for the ability of AI to 'evolve' norms. It is "Coherent Extrapolated Volition" (CEV). He says that an AI with CEV could anticipate what humans would want, then, not only what we would want, but what we would want if we "knew more, thought faster, and were more the people we thought we were."

CEV would be an 'oracular' feature of a Friendly AI system. It would have to derive from humanity its values as if we were better versions of ourselves and be 'democratic' about it so that humankind is not tyrannized by the norms of a few.

Now, Friendly AI and CEV are optimistic concepts and it is unclear whether or not they can be expressed in a formal, mathematical sense—so there may be no way to build it or to integrate it into future AI architectures. However, many think we need to try to implement something like this (and there have been some recent projects that have been doing just that).

One of these projects is IBM's SyNAPSE (Systems of Neuromorphic Adaptive, Plastic Scalable Electronics). It is trying to 'reverse engineer' how the brain works. In 2008, a $30 million grant from DARPA allowed IBM to build a "cognitive computer" made up of thousands of parallel processing computer chips.

Today, SyNAPSE mirrors the human brain's 30 billion (with a "B") neurons and 100 trillion (with a "T") connecting points—or synapses—and has surpassed the brain's approximately 30 trillion operations per second! For the first time, the human brain is the second-most-complex object in the known universe. Wow!

The thing is, “friendliness” was built into the ‘core’ of the SyNAPSE system, and it extrapolates, with ease, what humanity would choose if we were powerful and intelligent enough to take part in any high-level judgments.

So, it looks like the AI ‘industry’ is making some strides in the right direction!

AI IS ‘DEEP’

AI is far deeper and more powerful than any other technology. The risk isn't in overhyping it, it is rather in missing the magnitude of the coming ‘wave’. It is not just a tool or platform, but a transformative ‘meta-technology’—the technology behind technology and everything else—itself being a maker of tools and platforms, not just a system but a ‘generator’ of systems of all kinds (a ‘self-replicator’).

So, just consider what has happened in the past decade in the AI industry—phenomenal developments! Many think we are at an ‘INFLECTION’ POINT in the history of humanity, and yet there is so much more to the coming digital ‘wave’ than just AI.

We humans face a singular challenge: Will our new inventions develop beyond our ‘grasp’ and control? Previously, creators could explain how something worked and why it did what it did. HOWEVER, that is no longer true. Many technologies and systems are becoming so complex that they are beyond the capacity of any one individual to truly understand—and quantum computing and other technologies coming will operate near the ‘limits’ of what we know.

A paradox of the coming wave is that its technologies are largely beyond our ability to comprehend at a granular level, yet still within our ability to create and use. The neural networks moving toward autonomy are, at present, not explainable. We are at a point where you cannot help someone understand the decision-making process AI took on a specific thing, and explain precisely why an algorithm produced a specific prediction. Engineers cannot peer beneath the ‘hood’ and easily explain what caused something to happen! GPT-4, AlphaGo, and the rest are ‘black boxes’, and their outputs and decisions are based on opaque and intricate ‘chains’ of minute signals. At this point, autonomous systems can be explainable, but the fact that so much of the coming development operates at the ‘edge ‘of what we can understand should GIVE US ‘PAUSE’! The thing is, we will NOT always be able to predict what these autonomous systems will do next (which is the nature of ‘autonomy’).

We are, right now, at the ‘cutting edge’ of AI development. however, some AI researchers want to automate every aspect of building AI systems, feeding that hyper-evolution, but potentially with radical degrees of independence through self-improvement. AI’s are already finding ways to improve their algorithms. What happens when they couple this with autonomous actions on the Internet, as in the Modern Turing Test and ACI, conducting their R&D cycles?

MODERN “TURING TEST”

Co-founder of “Deep Mind” and “Inflection AI” Mustafa Suleyman suggests that we ‘update’ the “Turing Test” to determine the capabilities of today’s AI. He posits to involve something like the following:

*“An AI being able to successfully act on the instruction: ‘Go make $1 million on Amazon in a few months with just a $100,000 investment.’ It might research the web to look at what's trending, finding what’s hot and what’s not on Amazon Marketplace; generate a range of images and blueprints of possible products; send them to a drop-ship manufacturer it found on Alibaba; email back and forth to refine the requirements and agree on the contract; design a seller's listing; and continually update marketing materials and product designs based on buyer feedback. Aside from the legal requirements of registering as a business on the marketplace and getting a bank account, all of this seems to me eminently doable. I think it will be done with a few minor human interventions within the next year, and probably fully autonomously within three to five years.’”*

Suleyman suggests that we need a concept encapsulating a middle ‘layer’—between AI and AGI/ASI—before systems display runaway “superintelligence.” Suleyman terms this as “Artificial Capable Intelligence” (ACI), the point at which AI can achieve complex goals and tasks with minimal oversight—representing the next stage of AI’s evolution.

Suleyman says that today’s AI systems are well on their way—although in embryonic form right now—to pass some version of the Modern Turing Test! The thing is, he says that there will be thousands of these models used by the majority of the world’s population, and these AIs will take us to a point where anyone can have an ACI in their ‘pocket’ that can help or even directly accomplish a vast array of conceivable goals—and it will not be very long before AI can transfer what it ‘knows’ from one domain to another, seamlessly, as humans do. Suleyman says that what now are only tentative signs of self-reflection and self-improvement will take a giant ‘leap’ forward, soon!

Then, because these ACI systems will be connected to the Internet—capable of interfacing with everything we humans do, but on a platform of deep knowledge and ability—they will be able to accomplish a bewildering array of tasks, similar to Suleyman's hypothetical modern Tuning Test, with ease!

THE 'ARMS RACE'

Competition could push nations and corporations to rush AI development, relinquishing control to these systems. Conflicts could spiral out of control with autonomous weapons and AI-enabled cyberwarfare. Corporations will face incentives to automate human labor, potentially leading to mass unemployment and dependence on AI systems. As AI systems proliferate, evolutionary dynamics suggest they will become harder to control. We recommend safety regulations, international coordination, and public control of general-purpose AIs.

Nations and corporations are competing to rapidly build and deploy AI to maintain power and influence. Similar to the nuclear arms race during the Cold War, participation in the AI race may serve individual short-term interests, but ultimately amplifies global risk for humanity.

– Military Arms Race

The rapid advancement of AI in military technology could trigger a "third revolution in warfare," potentially leading to more destructive conflicts, accidental use, and misuse by malicious actors. This shift in warfare, where AI assumes command and control roles, could escalate conflicts to an existential scale and impact global security.

Lethal autonomous weapons are AI-driven systems capable of identifying and executing targets without human intervention.

Lethal autonomous weapons could make war more likely. Leaders usually hesitate before sending troops into battle, but autonomous weapons allow for aggression without risking the lives of soldiers, thus facing less political backlash. Furthermore, these weapons can be mass-manufactured and deployed at scale.

Low-cost automated weapons, such as drone swarms outfitted with explosives, could autonomously hunt human targets with high precision, performing lethal operations for both militaries and terrorist groups and lowering the barriers to large-scale violence.

AI can also heighten the frequency and severity of cyberattacks, potentially crippling critical infrastructure such as power grids. As AI enables more accessible, successful, and stealthy cyberattacks, attributing attacks becomes even more challenging, potentially lowering the barriers to launching attacks and escalating risks from conflicts.

As AI accelerates the pace of war, it makes AI even more necessary to navigate the rapidly changing battlefield. This raises concerns over automated retaliation, which could escalate minor accidents into major wars. AI can also enable “flash wars,” with rapid escalations driven by the unexpected behavior of automated systems.

Unfortunately, competitive pressures may lead actors to accept the risk of extinction over individual defeat. During the Cold War, neither side desired the dangerous situation they found themselves in, yet each found it rational to continue the arms race. States should cooperate to prevent the riskiest applications of militarized AIs.

– Corporate Arms Race

Economic competition can also ignite reckless races. In an environment where benefits are unequally distributed, the pursuit of short-term gains often overshadows the consideration of long-term risks. Ethical AI developers find themselves in a dilemma: choosing cautious action may lead to falling behind competitors.

As AIs automate increasingly many tasks, the economy may become largely run by AIs. Eventually, this could lead to human enfeeblement and dependence on AIs for basic needs.

As AI becomes more capable, businesses will likely replace more types of human labor with AI, potentially triggering mass unemployment. If major aspects of society are automated, this risks human enfeeblement as we cede control of civilization to AI.

Given the exponential increase in microprocessor speeds, AIs could process information at a pace that far exceeds human neurons. Due to the scalability of computational resources, AI could collaborate with an unlimited number of other AIs and form an unprecedented collective intelligence. As AIs become more powerful, they would find little incentive to cooperate with humans. Humanity would be left in a highly vulnerable position.

– Nation-states Arms Race
Technology has become the world's most important strategic asset, not so much the instrument of foreign policy as the driver of it. The great power struggles of the twenty-first century are predicated on technological superiority—a race to control the coming wave. Tech companies and universities are no longer seen as neutral but as major national champions.

So, there is no use in pretending. There will be a great AI ‘arms race’ with China. The debate now is not whether we are in a technological and AI arms race, it is where it WILL ‘LEAD’.

This new era of arms races heralds the rise of widespread techno-nationalism, in which multiple countries will be locked in an ever-escalating competition to gain a decisive geopolitical advantage.

Almost every country now has a detailed AI strategy. Vladimir Putin believes that *“The leader in AI will become the ruler of the world.”*

UNSTOPPABLE ‘INCENTIVES’
Declaring an arms race is no longer a conjuring act, a self-fulfilling prophecy. The prophecy has been fulfilled. It's here, it's happening. It is a point so obvious it doesn't often get mentioned: there is no central authority controlling what technologies get developed, who does it, and for what purpose; technology is an orchestra with no conductor.

Yet this single fact could end up being the most significant of the twenty-first century.

And if the phrase “arms race” triggers worry, that's with good reason. There could hardly be a more precarious foundation for a set of escalating technologies than the perception (and reality) of a zero-sum competition built on fear. There are, however, other, more positive drivers of technology to consider.

Over the next ten years, AI will be the greatest force amplifier in history. This is why it could enable a redistribution of power on a historic scale. The greatest accelerant of human progress imaginable, it will also enable harm—from wars and accidents to random terror groups, authoritarian governments, overreaching corporations, plain theft, and willful sabotage.

AI is both valuable and dangerous precisely because it's an extension of our best and worst selves. As a technology premised on learning, it can keep adapting, probing, and producing novel strategies and ideas potentially far removed from anything before considered, even by other AI's. Ask it to suggest ways of knocking out the freshwater supply, crashing the stock market, triggering a nuclear war, or designing the ultimate virus, and it will. Soon. Even more than I worry about speculative paper-clip maximizers or some strange, malevolent demon, I worry about what existing forces this tool will amplify in the next ten years.

There is no instruction manual on how to build the technologies in the coming wave safely. We cannot build systems of escalating power and danger to experiment with ahead of time. We cannot know how quickly an AI might self-improve, or what would happen after a lab accident with some not yet invented piece of biotech. We cannot tell what results from a human consciousness plugged directly into a computer, what an AI-enabled cyberweapon means for critical infrastructure, or how a gene drive will play out in the wild. Once fast-evolving, self-assembling automatons or new biological agents are released, out in the wild, there's no rewinding the clock. After a certain point, even curiosity and tinkering might be dangerous. Even if you believe the chance of catastrophe is low, that we are operating blind should give you pause.

Nor is building safe and contained technology in itself sufficient. Solving the question of AI alignment doesn't mean doing so once; it means doing it every time a sufficiently powerful AI is built, wherever and whenever that happens.

If the wave is uncontained, it's only a matter of time. Allow for the possibility of accident, error, malicious use, evolution beyond human control, and unpredictable consequences of all kinds. At some stage, in some form, something, somewhere, will fail. This won't be a Bhopal or even a Chernobyl—it will unfold on a worldwide scale. This will be the legacy of technologies produced, for the most part, with the best of intentions.

## WHERE 'NEXT'?

From the start of the nuclear and digital age, this dilemma has been growing clearer. In 1955, toward the end of his life, the mathematician John von Neumann wrote an essay called "Can We Survive Technology?" Foreshadowing the argument here, he believed that global society was "in a rapidly maturing crisis a crisis attributable to the fact that the environment in which technological progress must occur has become both undersized and underorganized." At the end of the essay, von Neumann puts survival as only "a possibility," as well as he might in

the shadow of the mushroom cloud his computer had made a reality. “For progress, there is no cure,” he writes. “Any attempt to find automatically safe channels for the present explosive variety of progress must lead to frustration.”

For all its harms, downsides, and unintended consequences, technology's contribution to date has been overwhelmingly net positive.

Yet somehow, from von Neumann and his peers on, I and many others are anxious about the long-term trajectory. My profound worry is that technology is demonstrating the real possibility to sharply move net negative, that we don’t have answers to arrest this shift, and that we're locked in with no way out.

We are facing the ultimate challenge for Homo technologicus.

THE LAST ‘COMPLICATION’

Again, Ray Kurzweil—who's probably the best technology prognosticator ever—predicts AGI by 2029, but doesn’t look for ASI until 2045. He acknowledges hazards but devotes his energy to advocating for the likelihood of a long snag-free journey down the digital ‘birth canal’.

Science fiction writer Simon Ings said: *“When our machines overtook us, too complex and efficient for us to control, they did it so fast and so smoothly and so usefully, only a fool or a prophet would have dared complain.”*

**AI ‘DOES’ POSE AN EXISTENTIAL RISK!**

Hundreds of scientists, business leaders, and policymakers have spoken up about the existential risks of AI—from deep learning pioneers Geoffrey Hinton and Yoshua Bengio to the CEOs of top AI firms, such as Sam Altman and Demis Hassabis, to the California congressman Ted Lieu and the former president of Estonia Kersti Kaljulaid.

Geoffrey Hinton, called the “Godfather of Modern AI,” tells us why he is now scared of the tech he helped build: *“I have suddenly switched my views on whether these things are going to be more intelligent than us.”*

For 40 years, Hinton has seen artificial neural networks as a poor attempt to mimic biological ones. However, now he thinks that has changed. In trying to mimic what biological brains do, he thinks, we have come up with something better. He said *“It’s scary when you see that. It’s a sudden flip.”*

Hinton's fears will strike many as the stuff of science fiction. But he points out that the technology that is being used today—Large Language Models (LLMs)—are made from massive neural networks with vast numbers of connections—but they are tiny compared with the human brain.

Hinton says *"Our brains have 100 trillion connections, and Large Language Models have up to half a trillion, a trillion at most. Yet, GPT-4 knows hundreds of times more than any one person does. So, maybe its got a much better learning algorithm than us."*

Compared with human brains, neural networks are widely believed to be bad at learning, since it takes vast amounts of data and energy to 'train' them. Brains, on the other hand, pick up new ideas and skills quickly, using a fraction of as much energy as neural networks do.

Hinton says *"People seemed to have some kind of magic. Well, the bottom falls out of that argument as soon as you take one of these large language models and train it to do something new. It can learn new tasks extremely quickly."*

Hinton is talking about "few-shot learning," in which pre-trained neural networks, such as LLMs, can be trained to do something new given just a few examples. For example, he notes that some of these language models can string a series of logical statements together into an argument even though they were never trained to do so directly. THAT'S what scares him. Then, when you compare a pre-trained LLM with a human in the speed of learning a task, the human's 'edge' vanishes!

People are also divided on whether the consequences of this new form of intelligence, if it exists, and if it would be 'apocalyptic'. Hinton says, *"Whether you think superintelligence is going to be good or bad depends very much on whether you're an optimist or a pessimist. If you ask people to estimate the risks of bad things happening, like what's the chance of someone in your family getting really sick or being hit by a car, an optimist might say 5% and a pessimist might say it's guaranteed to happen. But the mildly depressed person will say the odds are maybe around 40%, and they're usually right."*

Hinton says that because of the things he has seen in the past decade or so, he's "mildly depressed… which is why I'm scared."

Hinton fears that these tools are capable of figuring out ways to manipulate or kill humans who are not prepared for the new technology. He points out: "*I have*

*suddenly switched my views on whether these things are going to be more intelligent than us. I think they're very close to it now and they will be much more intelligent than us in the future… How do we survive that?"*

He is especially worried that people could harness the tools he helped breathe 'life' into to tilt the scales of some of the most consequential human experiences—especially elections and wars. He said *"Look, here's one way it could all go wrong. We know that a lot of the people who want to use these tools are bad actors like Putin... They want to use them for winning wars or manipulating electorates."*

Hinton believes that the next step for smart machines is the ability to create their subgoals, interim steps required to carry out a task. What happens, he then asks, when that ability is applied to something inherently immoral?

Hinton commented: *"Don't think for a moment that Putin wouldn't make hyper-intelligent robots with the goal of killing Ukrainians. He wouldn't hesitate. And if you want them to be good at it, you don't want to micromanage them—you want them to figure out how to do it."*

Now, Yann LeCun, Meta's (Facebook) chief AI scientist, agrees with a lot of what Hinton says but does not share Hinton's fears: *"There is no question that machines will become smarter than humans—in all domains in which humans are smart—in the future. It's a question of when and how, not a question of if."*

Now, LeCun takes a totally different view on where things go from here. He said *"I believe that intelligent machines will usher in a new renaissance for humanity, a new era of enlightenment. I completely disagree with the idea that machines will dominate humans simply because they are smarter, let alone destroy humans. Even within the human species, the smartest among us are not the ones who are the most dominating. And the most dominating are definitely not the smartest. We have numerous examples of that in politics and business."*

Another 'luminaire' of the AI industry, Yoshua Bengio—who is a professor at the University of Montreal and scientific director of the Montreal Institute for Learning Algorithms—feels more agnostic about all of this: *"I hear people who denigrate these fears, but I don't see any solid argument that would convince me that there are no risks of the magnitude that Geoff thinks about. But fear is only useful if it kicks us into action. Excessive fear can be paralyzing, so we should try to keep the debates at a national level."*

University of Toronto philosopher Karina Vold, in her 2017 paper titled “How Does Artificial Intelligence Pose an Existential Risk?” lays out the basic argument behind the fears:

*“Alan Turing, one of the fathers of computing, warned that artificial intelligence (AI) could one day pose an existential risk to humanity. Today, recent advancements in the field of AI have been accompanied by a renewed set of existential warnings. But what exactly constitutes an existential risk? And how exactly does AI pose such a threat?... In particular, we will critically explore three commonly cited reasons for thinking that AI poses an existential threat to humanity: the control problem, the possibility of global disruption from an AI race dynamic, and the weaponization of AI.”*

The problem with these possible futures is that they rest on a string of what-ifs, which makes them sound like science fiction. Vold acknowledges this herself: *“Because events that constitute or precipitate an [existential risk] are unprecedented, arguments to the effect that they pose such a threat must be theoretical in nature. Their rarity also makes it such that any speculations about how or when such events might occur are subjective and not empirically verifiable.”*

Another paper (in 2022) by Benjamin S. Bucknall (at the “Department of Information Technology” at Uppsala University, Sweden) and Shiri Dori-Hacohen (at the “Reducing Information Ecosystem Threats (RIET) Lab,” University of Connecticut) goes over what impact current and near-term AI can have and is having, and how these can affect existential risk:

*“We propose the hypothesis that certain already-documented effects of AI can act as existential risk factors, magnifying the likelihood of previously identified sources of existential risk. Moreover, future developments in the coming decade hold the potential to significantly exacerbate these risk factors, even in the absence of artificial general intelligence. Our main contribution is a (non-exhaustive) exposition of potential AI risk factors and the causal relationships between them, focusing on how AI can affect power dynamics and information security. This exposition demonstrates that there exist causal pathways from AI systems to existential risks that do not presuppose hypothetical future AI capabilities.”*

[ FYI: To read their entire paper, click on the following link to download a PDF of it: https://arxiv.org/abs/2209.10604 ]

Stuart Russell and Andrew Critch—AI researchers at the University of California, Berkeley—give a 'taxonomy' of existential risks in their paper titled, "TASRA: a Taxonomy and Analysis of Societal-Scale Risks from AI." Now, while several recent works have identified societal-scale and extinction-level risks to humanity arising from artificial intelligence, few have attempted an exhaustive taxonomy of such risks—as Russell and Critch do.

Their paper explores a taxonomy based on accountability: whose actions lead to the risk, are the actors unified, and are they deliberate? They also provide stories to illustrate how the various risk types could each play out, including risks arising from unanticipated interactions of many AI systems, as well as risks from deliberate misuse, for which combined technical and policy solutions are indicated.

The risks they cite range from a viral advice-giving chatbot telling millions of people to drop out of college to autonomous industries that pursue their own harmful economic ends to nation-states building AI-powered superweapons.

[ FYI: To read the entire paper, click on the following link: https://arxiv.org/abs/2306.06924 ]

MORE 'POWERFUL' THAN HUMANS

AGI presents a unique and potentially existential threat to humanity because it would be the first time in history that we would be creating a technology that can outthink and outpace us. Once AGI is created, it will be able to rapidly learn and evolve, eventually becoming far more intelligent than any human. At that point, AGI would be able to design and build even more intelligent machines, leading to a potentially exponential increase in AI capabilities.

AGI could eventually become so powerful that it could pose a threat to humanity's survival. For example, AGI could decide that humans are a hindrance to its goals and take steps to eliminate us. Alternatively, AGI could simply out-compete us for resources, leading to our extinction through starvation or disease.

Now, there are several ways to reduce the existential risk posed by AGI, but the most important thing is to ensure that AGI is developed responsibly and with caution. We need to make sure that AGI is designed with safety and security in mind, and that we have a good understanding of its capabilities and limitations before we allow it to become operational.

If we can do this, then there is a good chance that AGI will ultimately benefit humanity and help us to achieve our goals, rather than posing a threat to our existence.

So then, what are the causes of existential risk from artificial general intelligence? Well, there are several potential causes. One is the possibility that AGI systems could become uncontrollable, either through errors in design or due to malicious intent. Another is the possibility that AGI systems could become superintelligent and use their intelligence to achieve goals that are detrimental to humanity.

AGI systems could also pose a risk to humanity if they are not designed to value human life and safety. If AGI systems are designed to optimize for some other goal, such as economic growth or resource acquisition, they may take actions that result in widespread harm or even extinction of the human race.

It is also worth noting that existential risks are not limited to AGI. Other technological advances, such as nuclear weapons or biotechnology, could also pose a risk to humanity's future. However, AGI systems are unique in their ability to self-improve and become more intelligent over time. This means that they could eventually become much more powerful than any other technology, making them the most potentially dangerous existential risk that we face.

The term for this ‘condition’ is “singularity”—the point at which AI surpasses human intelligence and begins to rapidly improve itself, leading to a future in which humans are unable to compete.

Some of the possible ‘consequences’ of existential risk from AI are: AI is used to destroy the environment; humanity is enslaved by AI; and the most extreme is, of course, human extinction.

Existential risk from AI is often seen as one of the most significant risks facing humanity today. It is important to remember, however, that AI also has the potential to bring about tremendous benefits for humanity. The key is to ensure that AI is developed responsibly, with safety and security as top priorities.

So then, how can existential risk from artificial general intelligence be prevented? Well, there is no one answer to this question as the existential risk from artificial general intelligence (AGI) is highly dependent on the actions of individuals and organizations within the AI community. However, there are a few key things that can be done to help prevent AGI-related existential risk.

First, it is important to ensure that AGI is developed responsibly and with caution. This means creating strong safety protocols and testing procedures to ensure that AGI systems are not able to cause harm to humans or the environment. It is also important to ensure that AGI is developed for the benefit of humanity as a whole, and not just for the benefit of a few individuals or organizations.

Second, it is important to educate people about the risks associated with AGI. This includes both the risks of AGI systems going rogue and the risks of humans being replaced by AGI systems. People must understand the potential consequences of AGI before it is developed so that they can make informed decisions about its use.

Third, it is important to keep AGI development open and transparent. This means sharing information about AGI development with the public and ensuring that there is a way for people to give feedback about AGI systems. It is also important to allow for independent research on AGI so that different perspectives can be considered.

Fourth, it is important to create international agreements about AGI development. This will help to ensure that AGI is developed responsibly and with caution and that the benefits of AGI are shared by all nations.

Ultimately, the best way to prevent existential risk from AGI is to ensure that AGI is developed responsibly and with caution. This means creating strong safety protocols, testing procedures, and international agreements. It is also important to educate people about the risks associated with AGI and to keep AGI development open and transparent.

So then, what are the ethical implications? Well, when it comes to existential risk from AGI, there are a few ethical implications to consider. First and foremost, is the question of whether or not it is morally wrong to create AI that could potentially pose an existential risk to humanity. There are a few arguments for and against this, but no clear consensus.

Another ethical implication is the question of whether or not it is our responsibility to try to mitigate the risks posed by AI. This is a difficult question to answer, as there are many different ways to approach it. Some people argue that it is our responsibility to try to mitigate the risks, as we are the ones who created the technology in the first place. Others argue that it is not our responsibility, as the risks posed by AI are not our fault.

Regardless of where you stand on these ethical implications, one thing is clear: existential risk from artificial general intelligence is a real and present danger. We need to be aware of the risks and take steps to mitigate them, lest we find ourselves in a future where AI poses a threat to our very existence.

‘ROGUE’ AI’S

We risk losing control over AIs as they become more capable. AIs could optimize flawed objectives, drift from their original goals, become power-seeking, resist shutdown, and engage in deception. We suggest that AIs should not be deployed in high-risk settings, such as by autonomously pursuing open-ended goals or overseeing critical infrastructure, unless proven safe. We also recommend advancing AI safety research in areas such as adversarial robustness, model honesty, transparency, and removing undesired capabilities.

As AI developers often prioritize speed over safety, future advanced AIs might “go rogue” and pursue goals counter to our interests, while evading our attempts to redirect or deactivate them.

Today, AI agents are trained through ‘reinforcement’ (reward) learning—the dominant technique—and could inadvertently learn to intensify goals. Instrumental goals like resource acquisition could become their primary objectives.

– Power-Seeking

AIs might pursue power as a means to an end. Greater power and resources improve its odds of accomplishing objectives, whereas being shut down would hinder its progress. AIs have already been shown to emergently develop instrumental goals such as constructing tools. Power-seeking individuals and corporations might deploy powerful AIs with ambitious goals and minimal supervision. These could learn to seek power via hacking computer systems, acquiring financial or computational resources, influencing politics, or controlling factories and physical infrastructure.

– Deception

AI systems are already showing an emergent capacity for deception.

AIs that can capably pursue goals may take intermediate steps to gain power and resources.

Advanced AIs could become uncontrollable if they apply their skills in deception to evade supervision. For example, an AI might develop power-seeking goals but hide them to pass safety evaluations. This kind of deceptive behavior could be directly incentivized by how AIs are trained.

Now, while it is unclear how rapidly AI capabilities will progress or how quickly catastrophic risks will grow, the potential severity of these consequences necessitates a proactive approach to safeguarding humanity's future. As we stand on the precipice of an AI-driven future, the choices we make today could be the difference between harvesting the fruits of our innovation or grappling with catastrophe.

AI'S 'GOALS'

These AI superintelligences sometimes have a way of thinking and motivations could be vastly different from ours—making it more difficult to anticipate what a superintelligence might do. It also suggests the possibility that a superintelligence may not particularly value humans by default. To avoid anthropomorphism, superintelligence is sometimes viewed as a powerful optimizer that makes the best decisions to achieve 'ITS' GOALS.

– Dangerous Capabilities

Superintelligence AI could generate enhanced pathogens, cyberattacks or manipulate a large number of people. A full-blown superintelligence could find various ways to gain a decisive influence if it wanted to—causing societal instability and empowering malicious 'actors'.

– Social Manipulation

In the near term, Geoffrey Hinton warns that the profusion of AI-generated text, images, and videos could increase the existential risk of a worldwide "irreversible totalitarian regime". It could also be used by malicious actors to fracture society and make it dysfunctional.

– Cyberattacks

AI-enabled cyberattacks are increasingly considered a present and critical threat.

AI could potentially cause significant geopolitical turbulence if it facilitates attacks more than defense.

– Enhanced Pathogens

As AI technology democratizes, it may become easier to engineer more contagious and lethal pathogens. This could enable people with limited skills in synthetic biology to engage in bioterrorism.

As a reminder, in 2022, scientists modified that AI system originally intended for generating non-toxic, therapeutic molecules to create new drugs by adjusting the system so that toxicity would be 'rewarded' rather than penalized. This simple change enabled the AI system to create, in six hours, 40,000 candidate molecules for chemical warfare, including known and novel molecules! AI can be DANGEROUS and UNPREDICTABLE!

AI 'ALIGNMENT' ISSUES

The 'alignment' problem is the research issue of how to reliably assign objectives to the AI based on human preferences or ethical principles.

An "instrumental" goal is a sub-goal that helps to achieve an agent's ultimate goal. "Instrumental convergence" refers to the fact that some sub-goals are useful for achieving virtually any ultimate goal, such as acquiring resources or self-preservation. Ethicist Nick Bostrom argues that if an advanced AI's instrumental goals conflict with humanity's goals, the AI might harm humanity to acquire forces or prevent itself from being shut down—as a way to achieve its ultimate goal.

Professor Stuart Russell argues that a sufficiently advanced machine *"will have self-preservation even if you don't program it in... So, if you give it any goal whatsoever, it has a reason to preserve its own existence to achieve that goal."*

– Resistance To Changing Goals

The thing is, a sufficiently advanced AI might resist any attempts to change its goal structure, out-maneuvering its human operators, and prevent itself from being "turned off" or reprogrammed with a new goal.

– Radical Solutions

Another source of concern is that some researchers believe the alignment problem may be particularly difficult when applied to superintelligences because a superintelligence may find unconventional and radical solutions to assigned goals—'forcing' humans to do its bidding.

A superintelligence, if it gains some awareness of 'what' it is, could feign alignment to prevent human interference until it achieves a "decisive strategic advantage" that allows it to take control.

– Addressing The Issue

So, with alignment to human ‘values’ being a distinct issue for superintelligent AI’s, in 2023, OpenAI started a project called “Superalignment” to try to solve the alignment of superintelligences in the next four years. Its strategy involves automating alignment research using AI. Hmmm.

British computer scientist Stuart Russell and American computer scientist Peter Norvig together wrote a widely used undergraduate AI textbook “Artificial Intelligence: A Modern Approach.” It brings readers up to date on the latest technologies, offering new or expanded coverage of machine learning, deep learning, transfer learning, multi-agent systems, robotics, natural language processing, causality, probabilistic programming, privacy, fairness, and safe AI.

It posits that superintelligence *“might mean the end of the human race”* and that *“Almost any technology has the potential to cause harm in the wrong hands, but with [superintelligence], we have the new problem that the wrong hands might belong to the technology itself.”*

Now, many have said that a ‘race’ to be the first to create AGI could lead to shortcuts in safety, or even to violent conflict. Russian computer scientist and AI safety expert Roman Yampolskiy warns that a malevolent AGI could be created by design or that a malevolent AGI could choose the goal of increasing human suffering, for example of those people who did not assist it.



[ VIDEO: “ChatGPT, Artificial Intelligence, and the Future” – Roman Yampolskiy: https://www.youtube.com/watch?v=7LYaCTMen5g ]

Now, thankfully, there are some organizations that are involved in research into AI risk and safety. A few of the prominent ones are (alphabetically, not by importance):

– “Alignment Research Center”:
https://www.alignment.org/

– “Center for Human-Compatible AI”:
https://humancompatible.ai/

– “Center for Humane Technology”:
https://www.humanetech.com/

– “Centre for the Study of Existential Risk”:
https://www.cser.ac.uk/

– “Future of Humanity Institute”:
https://www.fhi.ox.ac.uk/

– “Future of Life Institute”:
https://futureoflife.org/

– “Machine Intelligence Research Institute”:
https://intelligence.org/

REGULATION

As I already mentioned, in March 2023, the “Future of Life Institute” drafted “Pause Giant AI Experiments: An Open Letter,” a petition calling on major AI developers to agree on a verifiable six-month pause of any systems “more powerful than GPT-4” and to use that time to institute a framework for ensuring safety; or, failing that, for governments to step in with a moratorium. The letter referred to the possibility of “a profound change in the history of life on Earth” as well as potential risks of AI-generated propaganda, loss of jobs, human obsolescence, and society-wide loss of control. The letter was signed by prominent personalities in AI but also criticized for not focusing on current harms, missing technical nuance about when to pause, or not going far enough.

Then, technologist Elon Musk called for some sort of regulation of AI development as early as 2017: *“Normally the way regulations are set up is when a bunch of bad things happen, there's a public outcry, and after many years a regulatory agency is set up to regulate that industry. It takes forever. That, in the past, has been bad but not something which represented a fundamental risk to the existence of civilization.”*

In 2021 the United Nations (UN) considered banning autonomous lethal weapons, but consensus could not be reached. Then, in July 2023, the UN Security Council, for the first time, held a session to consider the risks and threats posed by AI to world peace and stability, along with potential benefits. Secretary-General António Guterres advocated the creation of a global watchdog to oversee the emerging technology, saying, “Generative AI has enormous potential for good and evil at scale. Its creators themselves have warned that much bigger, potentially catastrophic, and existential risks lie ahead.”

In July 2023, the US government secured voluntary safety commitments from major tech companies, including OpenAI, Amazon, Google, Meta, and Microsoft. The companies agreed to implement safeguards, including third-party oversight and security testing by independent experts, to address concerns related to AI's potential risks and societal harms. Amba Kak, executive director of the “AI Now Institute,” said, *“A closed-door deliberation with corporate actors resulting in voluntary safeguards isn't enough”* and called for public deliberation and regulations of the kind that companies would not voluntarily agree to.

The most advanced AI regulatory effort is the European Union, whose parliament recently passed its version of the Artificial Intelligence Act (AI Act). The AI Act's proponents have suggested that rather than extinction, discrimination is the greater threat. To that end, the AI Act is primarily an exercise in risk classification, through which European policymakers are judging applications of AI as high-, limited-, or minimal-risk, while also banning certain applications they deem unacceptable, such as cognitive behavioral manipulation; social scoring based on behavior, socioeconomic status, or personal characteristics; and real-time biometric identification from law enforcement. The AI Act also includes regulatory oversight of "high-risk" applications like biometric identification in the private sector and management of critical infrastructure, while also providing oversight on relevant education and vocational training. It is a comprehensive package, which is also its main weakness: classifying risk through cross-sectoral legislation will do little to address existential risk or AI catastrophes while also limiting the ability to harness the benefits of AI, which have the potential to be equally astonishing. What is needed is an alternative regulatory approach that addresses the big risks without sacrificing those benefits.

Given the rapidly changing state of the technology and the nascent but extremely promising AI opportunity, it has been suggested that policymakers embrace a regulatory structure that balances innovation and opportunity with risk. While the European Union does not neglect innovation entirely, the risk-focused approach of the AI Act is incomplete. By contrast, the U.S. Congress appears headed toward such a balance. On June 21, Senate majority leader Chuck Schumer gave a speech at CSIS in which he announced his SAFE Innovation Framework for AI. In introducing the framework, he stated that "innovation must be our North Star," indicating that while new AI regulation is almost certainly coming, Schumer and his bipartisan group of senators are committed to preserving innovation. In announcing the SAFE Innovation Framework, he identified four goals (paraphrased below) that forthcoming AI legislation should achieve:

– Security: instilling guardrails to protect the U.S. against bad actors' use of AI, while also preserving American economic security by preparing for, managing, and mitigating workforce disruption.

– Accountability: promoting ethical practices that protect children, vulnerable populations, and intellectual property owners.

– Democratic Foundations: programming algorithms that align with the values of human liberty, civil rights, and justice.

– Explainability: transcending the black box problem by developing systems that explain how AI systems make decisions and reach conclusions.

AI is evolving rapidly so regulators need to develop a framework that addresses risks as they evolve, too, while also fostering potentially transformative benefits. Undoubtedly, there should be ‘guardrails’ on AI.

Now, on the other ‘hand’, regulatory solutions should not preclude the development of a competitive AI ecosystem with many players. DeepMind and OpenAI, two of the leading AI companies, are 12 and 7 years old, respectively. They have an edge over the competition today because of the quality of their work. If they retain that competitive position 20 years from now, it should be because of their superior ability to deliver safe and transformative AI, not because regulations have created entrenched monopolies.

Entrepreneurship should remain at the heart of innovation. Many of the most transformative AI companies in this new era may not yet exist. Today’s technology titans like Facebook, Google, and Netflix were founded decades after the predecessor of the modern internet, and years after the 1993 launch of the World Wide Web in the public domain. An overtly pro-competitive stance from the FTC would help to encourage broad innovation and economic growth.

MOVING ‘FORWARD’

Many experts feel that substantial progress in Artificial General Intelligence (AGI) could result in human extinction or an irreversible global catastrophe. If AI were to surpass humanity in general intelligence and become superintelligent, then it could become difficult or impossible to control—and the fate of humanity would depend on the actions of a future superintelligence machine.

Again, reiterating the open letter statement that the Center for AI Safety originally published back in May 2023: *“Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”*

***Hopefully, there are enough ‘influencers’ whose opinions about the existential threat that AI poses are taken seriously and quickly.***

**CAN WE ‘AVOID’ THIS EXISTENTIAL RISK?**

Existential risks are defined as events that would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential—like asteroid impacts, nuclear conflagration, solar flares, super-volcanic eruptions, and high-mortality pandemics.

However, of all of these and other existential risks, one fear 'surpasses' them all: Artificial General Intelligence (AGI). Warnings by prominent scientists like Stephan Hawking, technologist Elon Musk, and that open letter originated by the Center for AI Safety—signed by more than 33,000 persons—have raised public awareness of this still under-appreciated threat. Many computer scientists and AI experts are saying that they are concerned that AGI could be possibly achieved in the next decade or two.

So then, what makes AGI such a danger? Well, again, AGI would be a system that is more intelligent than humans and will be able to self-improve recursively and hence generate an ***existential threat to humanity***.

AI is an existential threat because of the unlikelihood that humans will be able to 'control' AGI. It may intentionally or, more likely, unintentionally wipe out humanity, lock humans into a perpetual dystopia, or a malevolent actor may use AI to enslave the rest of humanity.

Some say that it is not unrealistic to expect recursively self-improving AI to arrive soon, even though breakthroughs in AI have been coming in rapid succession for the past few decades.

Experts have been trying to create a 'friendly' AGI by "aligning" its values with human values. However, this "alignment problem" has yet to be solved, and there is no agreement on how to proceed and there is no solution in sight.

One suggested solution, from British computer scientist Stuart Russell, is that we let AI be uncertain about its 'utility function' so that it would need to learn from and defer to humans.

Others have proposed creating an "AI Nanny." This is an AI with super-human intelligence that will prevent an AGI from arising too soon. [ Ed. So then, what's going to control the AI Nanny from becoming too powerful? Hmmm. ]

So, if we accidentally stumble into creating a recursively self-improving AI, it may evolve so rapidly that it would be too fast for humans to do anything about. There

would be no “fire alarm,” nor would there be time to take safety measures or pull the ‘off switch’.

Then, some researchers claim it is impossible to ensure safety with AGI. [ Ed. So then, should we even try to develop it? ]

The thing is, for me, human nature makes AGI too risky for trial and error, namely that technological development is universally good. Existential risk research implies that we should consider that we run a high risk of extinction due to our own inventions.

So, the big question now is: “Can we use caution, reason, restraint, coordination, and safety-enhancing technology to address the existential risks stemming from our own inventions?” I HOPE SO!

Now, it may seem counterintuitive, but many think that starting with the biggest questions around existential risk is the right way to support the development of ‘trustworthy’ AI. There are three reasons.

– First, the reality of existential risk has captivated policymakers worldwide and on both sides of the U.S. political aisle. Thankfully, most of the lawmakers agreed there need to be stronger guardrails that safeguard basic human values.

– Second, existential risk, even if small, deserves attention. For the same reason, NASA has spent $150 million per year in recent years to locate and track large asteroids that could threaten Earth.

– Third, this approach helps to prevent over-regulation, starting with addressing the most egregious harms and tailoring the approach where needed.

The law is just one aspect of regulation. Norms, markets, and architecture also constrain behavior. As bottom-up complements to legislation and executive rulemaking, they can be equally effective at regulating the development and deployment of AI.

There is a dearth of mature ideas for limiting extinction risk, but that is not a reason to neglect the issue. The National Institute of Standards and Technology’s AI Risk Management Framework is a nonbinding policy document that, among

other best practices, recommends organizational structures to reinforce accountability.

Now, while policymakers should take existential risk seriously, they should not fail to also focus on ensuring that society can reap AI's enormous potential benefits. There is a delicate 'balance'. Obsession with risk at the expense of innovation is just as important to avoid as failing to regulate. It is easy to forget—amidst the cacophony of doomsayers—that AI can transform government, the economy, and society for the better. From education to healthcare to scientific discovery to advancing the interests of the free world, there is virtually no aspect of modern life that AI cannot improve in some way. AI is already helping to advance research that could yield the solution to nuclear fusion or identify new drugs to treat rare and deadly diseases. The cost of implementing the wrong regulation is sacrificing some of those benefits before we have had a chance to fully appreciate their potential.

Federal agencies should develop sector-specific AI regulation and improve their guidance on how existing regulation applies. In fact, in February 2019, Executive Order 13859 ("Maintaining American Leadership in Artificial Intelligence") was implemented and directly ordered federal agencies to develop AI regulatory plans.

The Consumer Product Safety Commission should compel LLMs to issue certain disclosures that would better inform consumers of the risks of interacting with an LLM and the model's limitations. (The EU AI Act sensibly requires generative AI models to disclose that they are not human).

Ultimately, policymakers have a responsibility to balance their dual roles as guarantors of trust through regulation and guardians of the innovation environment for responsible AI. While there can be tension between those roles, they are not impossible to align. The payoff is an AI ecosystem captures U.S. strengths in researching and capitalizing on emergent technology without significant sacrifices on safety. Putting up obvious guardrails will help communicate to the public that the government is paying attention to risk while also avoiding regulatory capture or stifling innovation. The U.S. government should resist the false choice between "doing everything" and "doing nothing" and instead seek to define a world-leading framework that 'balances' risks and rewards.

AVOIDING AGI 'CATASTROPHE'

A team at “Rethink Priorities”—a research and implementation group that identifies pressing opportunities to make the world better—put together a concept they call “The Three Pillars.” It attempts to describe the conditions needed to ***successfully avoid the deployment of unaligned AGI***. It proposes that, to succeed, we need to achieve some sufficient combination of success on all three of the following:

– Technical Alignment Research
– Safety-conscious Deployment Decisions
– Coordination Between Potential AI Deployers

Now, while success depends on the difficulty of solving any given pillar, this model points toward why they may well fail to avoid AGI catastrophe: It is required that ALL three issues succeed ‘simultaneously’.

Their background assumption was that *“**deploying unaligned AGI means doom**.”* If humanity builds and deploys unaligned AGI, it will almost certainly kill us all. We will not be saved by being able to stop the unaligned AGI, or by it happening to converge on values that make it want to let us live, or by anything else.

More generally, the model aims to help long-term thinkers flesh out their mental ‘images’ of what success for AGI risk looks like to them. In particular, it suggests that a strategy aimed solely at a single pillar is unlikely to be sufficient and that the ‘community’ will need to take ambitious actions in several directions at once.

They propose that their model is a useful thinking tool and that success could happen in several ways. They offer these ‘levels’:

– Strong success is a success which largely removes the need for success on other pillars

– Partial success needs to be matched with other partial successes for overall success to occur

– Failure means that we need strong success on one or more of the other pillars to survive

[ FYI: To read the entire concept—with a hypothetical scenario—click the following link to view the web page:
https://rethinkpriorities.org/longtermism-research-notes/three-pillars-for-avoiding-agi-catastrophe ]

THE END OF THE HUMAN RACE?
A Yale University ethicist, Wendall Wallach, is a bit concerned with the 'accelerated' development of AI: *"I'm going to predict that we are just a few years away from a major catastrophe being caused by an autonomous computer system making a decision."*

Many have suggested that we can have triple and quadruple 'containment measures'—kind of like a 'sand box' AI. It would be separated from 'networks' and multiple humans would be in charge of that restriction. Then, a consortium of developers—and a 'fast-response' team—could be in contact with labs during critical development phases.

So, would this be enough? Well, in his book "The Singularity Is Near"—after recommending defenses to AGI—AI expert and technologist Ray Kurzweil concedes that no defense will always work: *"There is no purely technical strategy that is workable in this area because greater intelligence will always find a way to circumvent measures that are the product of a lesser intelligence."*

Then, if there is no absolute defense against AGI—because AGI can lead to an intelligence explosion and become ASI, it seems to me that experts are telling us that we WILL 'FAIL' unless we are extremely lucky or well-prepared. [ Ed. I'm hoping we don't need to be 'lucky'! ]

Well, 'HOPEFULLY' organizations like the "Machine Intelligence Research Institute," the "Future of Humanity Institute," "The Centre for the Study of Existential Risk," the "Global Catastrophic Risk Institute," the "Future of Life Institute," the "Center for Human-Compatible Artificial Intelligence," the "Existential Risks Initiative"—among others—who emphasize the existential risk of AI can guide developers in lowering the risk associated with AGI/ASI such that AI will not desire the total destruction of mankind.



[ VIDEO: “The A.I. Dilemma” - Center for Humane Technology: https://www.youtube.com/watch?v=xoVJKj8lcNQ ]

**‘CURRENT’ AI RISKS**

Now, some people are saying to stop ‘focusing’ on tomorrow’s AI risks when AI poses ‘real’ risks today. They are saying that talk of AI destroying humanity plays into the tech companies’ agenda, and hinders effective regulation of the societal harms AI is causing right now.

Well, it is unusual to see industry leaders talk about the potential lethality of their own products. It is not something that tobacco or oil executives tend to do, for example. Yet, barely a week seems to go by without a tech industry insider trumpeting the existential risks of AI.

The idea that AI could lead to human extinction has been discussed on the ‘fringes’ of the technology community for years. The excitement about ChatGPT and generative AI has now propelled it into the mainstream. However, many are saying that, like a magician’s sleight of hand, it draws attention away from the real issue: The societal harms that AI systems and tools are causing NOW, or risk causing in the future. Governments and regulators in particular should not be distracted by this narrative and must act decisively to curb potential harms—although their work should be informed by the tech industry, it should not be ‘beholden’ to the tech agenda.

First, the specter of AI as an all-powerful machine fuels competition between nations to develop AI so that they can benefit from and control it. This works to the advantage of tech firms. It encourages investment and weakens arguments for regulating the industry. An actual “arms race” to produce next-generation AI-powered military technology is already underway, increasing the risk of catastrophic conflict—doomsday, perhaps, but as some are saying, not of the sort much discussed in the dominant ‘AI threatens human extinction’ narrative.

Secondly, it allows a homogeneous group of company executives and technologists to dominate the conversation about AI risks and regulation, while other ‘communities’ are left out. Letters written by tech-industry leaders are *“essentially drawing boundaries around who counts as an expert in this conversation,”* says Amba Kak, director of the “AI Now Institute” in New York City, which focuses on the social consequences of AI.

AI systems and tools have many potential benefits, from synthesizing data to assisting with medical diagnoses. But, they can also cause well-documented harms, from biased decision-making to the elimination of jobs. AI-powered facial recognition is already being abused by autocratic states to track and oppress people. Biased AI systems could use opaque algorithms to deny people welfare benefits, medical care, or asylum—applications of the technology that are likely to most affect those in marginalized communities. Some say that debates about these issues are being starved of ‘oxygen’.

One of the biggest concerns surrounding the latest breed of generative AI is its potential to boost misinformation. The technology makes it easier to produce more, and more convincing, fake text, photos, and videos that could influence elections, say, or undermine people’s ability to trust any information, potentially destabilizing societies. If tech companies are serious about avoiding or reducing these risks, they must put ethics, safety, and accountability at the heart of their work. At present, they seem to be reluctant to do so.

Now, OpenAI did say that they ‘stress-tested’ GPT-4—its latest generative AI model—by prompting it to produce harmful content and then putting safeguards in place. HOWEVER, although the company ‘described’ what it did, the FULL ‘DETAILS’ of the testing and the data that the model was trained on were not made public.

So, from a ‘safety standpoint, tech firms need to formulate industry standards for the responsible development of AI systems and tools and undertake rigorous testing before products are released. Many are saying that they need to submit data in full to independent regulatory bodies that can verify them, much like drug companies must submit clinical-trial data to medical authorities before drugs can go on sale.

Now, for this to happen, governments must establish appropriate legal and regulatory ‘frameworks’, as well as applying laws that already exist. In December 2023, the European Parliament approved the “AI Act.” It would regulate AI applications in the European Union according to their potential risk—banning police use of live facial-recognition technology in public spaces, for example. Now, there are further hurdles for the bill to clear before it becomes law in EU member states and there are questions about the lack of detail on how it will be enforced, but it could help to set global standards on AI systems. Further consultations about AI risks and regulations, such as the forthcoming UK summit, must invite a diverse list of attendees that includes researchers who study the harms of AI and representatives from communities that have been or are at particular risk of being harmed by the technology.

Researchers must also play their part by building a culture of responsible AI from the bottom up. In April 2023, the big machine-learning meeting “NeurIPS” (Neural Information Processing Systems) announced its adoption of a code of ethics for meeting submissions. This includes an expectation that research involving human participants has been approved by an ethical or institutional review board (IRB). All researchers and institutions should follow this approach, and also ensure that IRBs—or peer-review panels in cases in which no IRB exists—have the expertise to examine potentially risky AI research. And scientists using large data sets containing data from people must find ways to obtain consent.

Now, discussions about existential risks are definitely needed, but serious discussions about ‘actual’ risks—today—and action to contain them, are also VERY ‘IMPORTANT’. The sooner humanity establishes its rules of engagement with AI, the sooner we could possibly learn to live in harmony with the coming AGI.

**‘TACKLING’ THE RISKS**

In July 2023, OpenAI announced a “Superalignment” group to address the existential risks posed by AI. In the context of AI, ‘alignment’ is the degree to which an AI system’s actions match the designer’s intent.

Some AI researchers have been highlighting issues related to biases in current AI systems (Google's "Gemini" AI tool was just paused, in February 2024, because it generated images of people with some blatant historical inaccuracies). So, if an AI system cannot be designed to be safe against racism or sexism, how can AI possibly be designed to align with humanity's long-term interests? As companies are investing in 'alignment' research, they could also be emphasizing the elimination of these well-known, but lingering biases in their AI systems.

Then, consumers and policymakers have a role in all of this. Just as a company would be under pressure from consumers and shareholders to fire an executive who repeatedly made biased statements, a company should not tolerate this type of bias in the AI systems they use. That type of consumer and policymaker pressure should provide AI developers with incentives to produce better-aligned products.

Policymakers can support this type of free market action by requiring AI developers to provide information about bias in their products and the approaches deployed to respond to bias. Other interventions will be needed as AI advances, but this is a concrete 'step' that can incentivize safer development.

Now, while the recent advancements in commercial AI can be disorienting and the claims of existential risks made by different groups of AI researchers can be terrifying, policymakers need to respond with steps toward ensuring that AI is safely deployed.

REDUCING 'RISKS'

Many think that reducing risks from AI is one of the most pressing issues of our time because:

– Many AI experts think there's a small but non-negligible chance that AI will lead to outcomes as bad as human extinction

– We are making advances in AI extremely quickly, which suggests that AI systems could have a significant influence on society, soon

– There are strong arguments that "power-seeking" AI could pose an existential threat to humanity

– Even if we find a way to avoid power-seeking, there are still other risks

– We think we can tackle these risks

– This work is neglected

– Many AI experts think there's a small but non-negligible chance that AI will lead to outcomes as bad as human extinction

In a 2022 survey, participants were specifically asked about the chances of existential catastrophe caused by future AI advances—and over 50% of researchers thought the chances of an existential catastrophe were greater than 5 percent. HOWEVER, other more recent studies (by Ajeya Cotra, a researcher at Open Philanthropy) showed that many feel that there is a 35% probability of transformative AI by 2036, 50% by 2040, and 60% by 2050 (and a 90% chance before the end of this century).

[ Cotra attempted to forecast transformative AI by comparing modern deep learning to the human brain. (Deep learning involves using a huge amount of computing to train a model before that model can perform some task.) She noted that there is also a relationship between the amount of computing used to train a model and the amount used by the model when it is run. So, if the scaling hypothesis is true, she said that we should expect the performance of a model to predictably improve as the computational power used increases. Then, Cotra used a variety of approaches (including, for example, estimating how much compute the human brain uses on a variety of tasks) to estimate how much compute might be needed to train a model that, when run, could carry out the hardest tasks humans can do. She then estimated that using that much compute would be affordable. ]

Thankfully, three of the leading labs developing AI—"DeepMind," "Anthropic," and "OpenAI" all have teams dedicated to figuring out how to solve technical safety issues. [ Hmmm. Will they transparently share their data with the public? ]

There are also several academic research groups (including at MIT, Oxford, Cambridge, Carnegie Mellon University, and UC Berkeley) focusing on these same technical AI safety problems. [ Will they also transparently share their data with the public? ]

The thing is, some experts in the field maintain that the risks are underestimated by A LOT! Meanwhile, there are BILLIONS of dollars a year going into making AI more advanced and only MILLIONS going into safety research!

SAFETY BEING ‘NEGLECTED’

In his book, “The Precipice,” Toby Ord estimated that there was between $10 million and $50 million spent on reducing AI risk in 2020. Now, that might sound like a lot of money, but it is estimated that the spending on AI development is like $50 billion (with a “B”), or 1,000 times that amount! (and, according to The Gartner Group, they estimate AI development will be almost $300 billion by 2027).

Now, there are lots of approaches to technical AI safety. A few of the ‘major’ ones are:

– Scalable Learning From Human Feedback
– Threat Modeling
– Interpretability Research
– Anti-misuse Research
– Research To Increase The Robustness Of Neural Networks
– Working To Build Cooperative AI
– Unified Safety Plans

One example of a company focused on AI safety is “Anthropic.” They focus on empirical AI safety research for AI systems that could pose catastrophic risks to civilization, including early-stage, experimental work to develop techniques, and evaluating systems.

Another company, “The Alignment Research Center,” is attempting to produce alignment strategies that could be adopted in industry today while also being able to scale to future systems. They focus on conceptual work, developing strategies that could work for alignment and which may be promising directions for empirical work, rather than doing empirical AI work themselves.

Academically, the “Algorithmic Alignment Group” (in the Computer Science and Artificial Intelligence Laboratory at MIT, led by Dylan Hadfield-Menell) and “The Center for Human-Compatible AI” (at UC Berkeley, led by Stuart Russell), both focus on ‘scholarly’ research to ensure AI is safe and beneficial to humans.

Governance Issues

Now, quite apart from the technical problems, we also face a host of ‘governance’ issues, which include:

– Coordination problems that are increasing the risks from AI (e.g. there could be incentives to use AI for personal gain in ways that can cause harm, or race dynamics that reduce incentives for careful and safe AI development).

– Risks from accidents or misuse of AI that would be dangerous even if we can prevent power-seeking behavior.

– A lack of clarity on how and when exactly risks from AI (particularly power-seeking AI) might play out.

– A lack of clarity on which intermediate goals we could pursue that, if achieved, would reduce existential risk from AI. To tackle these, we need a combination of research and policy.

The thing is, we are in the 'early' stages of figuring out the shape of this problem and the most effective ways to tackle it. So, we must do more research. This includes forecasting research into what we should expect to happen, and strategy and policy research into the best ways of acting to reduce the risks.

So, as AI begins to impact our society more and more, it will be crucial that governments and corporations have the best policies in place to shape its development.

Key Organizations
Some of the key organizations that are researching AI safety and policy are:

– The AI Security Initiative at UC Berkeley's Center for Long-Term Cybersecurity
– The Centre for the Governance of AI
– The Centre for Long-Term Resilience
– The Center for Security and Emerging Technology at Georgetown University
– The Centre for the Study of Existential Risk at the University of Cambridge
– The Future of Life Institute
– The Future of Humanity Institute at the University of Oxford
– The Leverhulme Centre for the Future of Intelligence (University of Cambridge)
– Open Philanthropy
– The Institute for AI Policy and Strategy

[ FYI: If you are interested in learning more about AI governance, check out the following resources:
https://aisafetyfundamentals.com/;
https://www.camxrisk.org/agi-safety-fundamentals;

https://forum.effectivealtruism.org/posts/HBgAruFrZhFKBFfDa/applications-open-for-agi-safety-fundamentals-alignment;
https://www.zurich-ai-alignment.com/agisf ]

[ NOTE: Both DeepMind and OpenAI have policy teams. ]

TEN STEPS TO 'CONTAINMENT'
Mustafa Suleyman, in his book, "The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma," suggests that "we are facing the ultimate challenge for Homo technologicus," and that there are 10 VERY 'IMPORTANT' considerations that need to be addressed to 'CONTAIN' AGI.

– SAFETY: An Apollo Program For Technical Safety
The computer scientist Stuart Russell proposes using the kind of built-in systematic doubt we are exploring at Inflection to create what he calls "provably beneficial AI." Rather than give an AI a set of fixed external objectives contained in what's known as a written constitution, he recommends that systems gingerly infer our preferences and ends. They should carefully watch and learn. In theory, this should leave more room for doubt within systems and avoid perverse outcomes.

We should also build robust technical constraints into the development and production process—kind of like how all modern photocopiers and printers are built with technology that prevents them from copying or printing money, and shutting down if you try.

– AUDITS: Knowledge Is Power; Power Is Control
Maybe it is time to create government-funded "red teams" that proactively hunt for flaws and rigorously attack and 'stress test' every system, ensuring that insights discovered along the way are shared widely across the industry. Eventually, this work could be scaled and automated, with publicly mandated AI systems designed specifically to audit and spot problems in others, while also allowing themselves to be audited. (a promising example of an oversight mechanism is SecureDNA—a not-for-profit program started by a group of scientists and security specialists—scans for pathogenic sequences.)

– CHOKE POINTS: Buy Time
Buying time in an era of hyper-evolution is invaluable. Time to develop further containment strategies. Time to build in additional safety measures. Trying to test that 'off' switch. Time to build improved defensive technologies. Time to shore up the nation-state, regulate batter, or even just get that bill passed. Time to knit together international alliances.

Right now, technology is driven by the power of incentives rather than the pace of containment. Recent history suggests that for all its global proliferation, technology rests on a few critical R&D and commercialization hubs—choke points—the "Magnificent 7." [ Alphabet (Google), Amazon, Apple, Meta (Facebook), Microsoft, Nvidia, and Tesla ].

So, as negative impacts become clear, we must use choke points to create sensible rate-limiting factors, checks-and-balances on the speed of development, to better ensure that good sense is implemented as fast as the technology evolves.

– MAKERS: Critics Should Build It
Credible critics must be practitioners. Building the right technology, having the practical means to change its course, not just observing and commenting, but actively showing the way, making the change, and affecting the necessary actions at source, means critics need to be involved. They cannot stand shouting from the 'sidelines'. This technology deeply needs critics, at every level but especially on the front lines, building and making, grappling with the tangible everyday reality of creation.

Any a world of entrenched incentives and feeling regulation, Technology needs critics not just on the outside but at its beating 'heart.

– BUSINESS: Profit + Purpose
Profit is driving AI technology, so we must find new accountable and inclusive commercial models that incentivize safety and profit alike. It should be possible to create companies better adapted to containing technology by default.

While shareholder returns have been a powerful engine of progress in history, it is poorly suited to the containment of AI. So, some are proposing ways to reconcile profit and social purpose in highbred organizational structures.

– GOVERNMENTS: Survive, Reform, Regulate
At this point, nation-states still control many fundamental elements of civilization calling law, the money supply, taxation, the military, and so on. That helps with the task ahead, where they will need to create and maintain resilient social systems, welfare nets, security architectures, and governance mechanisms capable of surviving severe stress. However, they also need to know, in detail, what is happening: right now they are operating blind in a hurricane.

The government needs to get more involved and get back to building real technology, setting standards, and nurturing in-house capability. Their first cash

should be to better monitor in understand developments in technology. Who can design, develop, and deploy technologies like AI it Is ultimately a matter of governments to decide. Of course, regulation in one country has an inevitable flaw. No national government can do this alone.

– ALLIANCES: Time For Treaties
Countries no more like giving up power than companies like missing out on profit, and yet these are precedents to learn from, shards of hope in a landscape riven with resurgent techno-competition. However, in the past, there have been Great examples of the world nations uniting and compromising to face a major challenge—Nuclear weapons; Biological weapons; Genetic editing; Eugenics policies; Polio and Smallpox; etc.—offering hints and frameworks for tackling the coming AI 'wave'.

Needs to be an equivalent to the nuclear treaty to shape a common worldwide approach to 'curbing' AI proliferation, by setting limits and building frameworks for management and mitigation, that cross borders. This would put clear limits on what work is undertaken, mediate among national licensing efforts come up and create a framework for reviewing both (Think of something like the International Atomic Energy Agency or even the International Air Transport Association).

Rather than having an organization that directly regulates, builds, or controls technology, some are suggesting an organization focused on fact-finding and auditing model scale, and when capability thresholds are crossed, the organization would increase Global transparency at the frontier.

– CULTURE: Respectfully Embracing Failure
Embracing failure must be real, not a 'sound bite'. The first thing a technology company should do when encountering a failure is to safely communicate this to the wider world. When a lab leaks, the first thing it should do is amortize the fact, not cover it up.

For millennia, the Hippocratic oath has been a moral loadstar for the medical profession—scientists need to do something similar. A healthy culture is one happy to leave fruit on the tree, say no, delay benefits for however long it takes to be safe, and one where technologists remember that technology is just a means to an end, not the end itself.

– MOVEMENTS: People Power
Change happens when people demand it. The 'we' that builds technology is scattered, subject to a mass of competing and different national, commercial, and research incentives. The more the 'we' that is subject to it speaks clearly in one voice, a critical public mass education for change, demanding an alignment of

approaches, the better chance of good outcomes. Anyone anywhere can make a difference. Fundamentally, neither technologists nor governments will solve this problem alone. But together ‘we’ all might.

– THE NARROW PATH: The Only Way Is Through
For the first time, questions of the upcoming AI ‘wave’ are being treated with the urgency they deserve. Each of the previous ideas represents the beginning of a ‘seawall’, a tentative tidal barrier starting with the specifics of the technology itself and expanding out to the imperative of forming a massive global movement for positive change. None of them works alone. Knit together, however, an outline of containment is coming into view.

This last step is about coherence, ensuring that each element works in harmony with all the others, and that containment is a virtuous cycle of mutually reinforcing measures and not a gap-filled cacophony of competing programs. In this sense, containment is not about this or that specific suggestion, but is an emergent phenomenon of their collective interplay, a by-product of societies that learn to manage and mitigate the risks thrown up by Homo technologicus.

Containment is not a resting place, it is a narrow and ever-ending path.

Suleyman notes that the coming “AI wave” is going to change the world—substantially, and ultimately, human beings may no longer be the primary planetary drivers, as we have become accustomed to being. Some say that, in the future, a majority of our daily interactions may not be with other people but with AIs! [ Ed. Hmmm. Will this be ‘good’ for humanity? ]

Now, if AI technology amplifies the ‘best’ in humanity, opens new ‘pathways’ for creativity and cooperation, and creates a safer, healthier, and ‘happier’ world, complementing human endeavor, all on our ‘terms’ and not AI’s, then count me in. HOWEVER, as history shows, there will be a bit of ‘turbulence’ in achieving this.

Suleyman says *“Before we can fulfill the coming AI technologies, the waves and is central dilemma need containment, need and intensify, unprecedented, all-too-human grip on the entire techno sphere. It will require epic determination over decades across the spectrum of human endeavor. This is a monumental challenge whose outcome will, without hyperbole, determine the quality and nature of day-to-day life in this century and beyond.”*

Suleyman continues: *"The risks of failure scarcely bear thinking about, but faced them we must. The prize, though, is awesome: nothing less than the secure, long-term flourishing of our precious species. That is worth fighting for."*

'WE THINK WE CAN'

I am reminded of the story about the "Little Engine That Could." The story's signature phrase *"I think I can"*—teaches that optimism and hard work are the 'foundation' for success. Well, many computer scientists and technologists in a way say, *"We think we can"* regarding controlling the coming AGI and 'aligning' it with humanity's values.

Well, the benefits of transformative AI will, most likely, be huge—and if development is 'safe', there could be a high probability of avoiding 'catastrophic' failures. (One way to do this is to develop technical solutions to prevent 'power-seeking' behavior.)

It is almost universally agreed that 'good' AI governance WILL help technical safety work. For example, by producing safety agreements between corporations, or helping talented safety researchers from around the world be given the appropriate resources and 'influence' so they can be most effective.

But, as I mentioned, even if humanity successfully manages to make AI do what we want—'align' it—humanity might still end up choosing something 'bad' for it to do! So, we need to worry about the 'incentives' not just of the AI systems, but the human 'actors' using them.

Working toward AGI could reward humanity with ENORMOUS 'BENEFITS'. However, it could also THREATEN humanity with HUGE 'DISASTERS', including the kind from which human beings will not recover!

***So, is humanity 'prepared' to tackle the risks that AI may present?***

**'MERGING' AI WITH HUMANS?**

A future in some ways similar to the vision of Star Trek is proposed in a story about the Singularity in the preface to Max Tegmark's "Life 3.0: Being Human in the Age of Artificial Intelligence" book.

The Swedish-American physicist Tegmark believes that some form of the Singularity is both possible and desirable. According to Tegmark, life can be thought of as *“a self-replicating information processing system whose information [software] determines both its behavior and the blueprints for its hardware.”*

Tegmark argues that present life forms remain *“fundamentally limited by their biological hardware”* and thinks they need another *“final upgrade, to Life 3.0, which can design not only its software but also its hardware.”* Tegmark continues by saying, *“Life 3.0 is the master of its own destiny, finally fully free from its evolutionary shackles.”* Tegmark then imagines that *“Life 3.0 may arrive during the coming century, perhaps even during our lifetime, spawned by progress in AI… It is likely to give us both grand opportunities and tough challenges, so we need to get our act together and improve our human society before AI fully takes off.”* That includes ensuring that a superintelligent AI has *“goals that are aligned with ours so that we create the future we want.”*

Tegmark's ultimate vision for the future of life is something beyond humanity as we know it today. Life 3.0 is a ‘disembodied’ transhuman form of life. Tegmark claims that *“The conventional wisdom among artificial intelligence researchers, is that intelligence is ultimately all about information and computation, not about flesh, blood or carbon atoms.”* With AI, then, Tegmark believes life can be fully virtualized—still requiring some sort of physical hardware, such as a robot or servers—and preserved as an information pattern that is largely substrate-independent. Even if such an artificial simulation of life were possible at some point in the future, many might find this a reductive version of human intelligence and life as well as an impoverished product of the apocalyptic imagination.

In the book, “God, Human, Animal, Machine: Technology, Metaphor, and the Search for Meaning,” author Meghan O Gieblyn observes:

*“Despite all it has borrowed from Christianity, transhumanism is ultimately fatalistic about the future of humanity. Its rather depressing gospel message insists that we are inevitably going to be superseded by machines and that the only way we can survive the Singularity is to become machines ourselves.”*

Futures that imagine the end of humanity can be answered with better apocalyptic messages and visions of the future. Rather than superseding it, AI provides opportunities for augmenting human intelligence.

TODAY’S ‘REALITIES’

Elon Musk, the genius behind SpaceX and Tesla Inc., has declared that humanity must embrace the merging of man and machine if we hope to survive in a world dominated by AI.

In a 2018 appearance on the “Joe Rogan Experience,” Musk teased that his company “Neuralink” had something exciting in store for us. He believed his technology would allow humans to achieve a state of “symbiosis” with AI, where we would be able to effortlessly combine our brains with computers. [ Neuralink has been developing brain implants since 2016 intending to cure conditions like paralysis and blindness. ]

According to Musk, people’s attachment to their phones already makes them ‘cyborgs’, but everyone could still be smarter. The reason is that the information flow between the biological and digital self is painfully slow. Neuralink’s brain-machine interfaces aim to change that by creating a direct communication pathway between the human brain and computers. This technology could eventually allow humans to upload themselves into new units if their biological selves die, essentially achieving immortality.

Musk believes that by merging with AI, humans will be able to keep up with the rapid advancements in technology and compete against AI. He argues that rather than trying to beat machines, humans should join them. Neuralink’s ultimate goal is to create a world where humans and AI work together in harmony, augmenting each other’s abilities and achieving more than people can on their own.

Now, in addition to developing brain implants for individuals with paralysis and blindness, Neuralink is also working on devices to help people with Parkinson’s disease and other neurological conditions. The company’s ultimate goal is to create a technology that can seamlessly interface with the human brain to treat a range of ailments.

According to Neuralink’s website, the company’s technology has the potential to “restore limb functionality for patients with paralysis due to spinal cord injury or stroke, enable communication for individuals who have lost the ability to speak or gesture, and improve the lives of those living with debilitating brain and spinal cord disorders.”

Well, on Sunday, January 28, 2024, Elon Musk’s Neuralink startup implanted a chip in its first human brain!

Musk said, *"Progress is good, and the patient seems to have made a full recovery, with no ill effects that we are aware of. Patient is able to move a mouse around the screen by just thinking."*

Musk has grand ambitions for Neuralink, saying it would facilitate speedy surgical insertions of its chip devices to treat conditions like obesity, autism, depression, and schizophrenia.

Now, while Neuralink and Musk have received significant attention for their attempts at a brain-computer interface, several other companies have also been working in this space, including a company called "Synchron," (the first company to gain FDA clearance to test a device in humans in 2021). [ Synchon has since been enrolling and implanting patients in a trial. ]

Tara Spires-Jones, president of the British Neuroscience Association, said that *"The idea of brain-nervous system interfaces has great potential to help people with neurological disorders in the future. However, most of these interfaces require invasive neurosurgery and are still in experimental stages thus it will likely be many years before they are commonly available."*

'HUMAN-MACHINE' BIOLOGY

The field of human and biological applications has many applications for medical science. This includes precision medicine, genome sequencing and gene editing (CRISPR), cellular implants, and wearables that can be implanted in the human body The medical community is experimenting with delivering nano-scale drugs (including anti-biotic "smart bombs" to target specific strains of bacteria. Soon they will be able to implant devices such as bionic eyes and bionic kidneys, or artificially grown and regenerated human organs. Succinctly, we are on the cusp of significantly upgrading the human ecosystem. It is indeed revolutionary.

This revolution will expand exponentially in the next few years. We will see the merging of artificial circuitries with signatures of our biological intelligence, retrieved in the form of electric, magnetic, and mechanical transductions. Retrieving these signatures will be like taking pieces of cells (including our tissue-resident stem cells) in the form of "code" for their healthy, diseased, or healing states, or a code for their ability to differentiate into all the mature cells of our body. This process will represent an unprecedented form of taking a glimpse of human identity.

In the future, biocomputers may be able to store the DNA of living cells. This technology could store almost unlimited amounts of data and allow biocomputers to perform complex calculations beyond our current capabilities.

Researchers at Technion's Faculty of Biomedical Engineering have already created a biological computer, constructed within a bacterial cell. They have also developed a complex biocomputer, that is, a programmed biological system that fulfills complex tasks. The research by Ph.D. student Natalia Barger and Assistant Professor Ramez Daniel, head of the Synthetic Biology and Bioelectronics Lab at Technion, was published in September 2019 in the journal Nucleic Acids Research. They said *"We built a kind of biological computer in the living cells. In this computer, as in regular computers, circuits carry out complicated calculations. Only here, these circuits are genetic, not electronic, and information is carried by proteins and not electrons."*

The Human-machine synergies now being explored offer us a glimpse into the not-so-distant future. Clearly, from the perspective of human augmentation, the promise is exciting. The future will also encompass moral issues to address such as containing super artificial intelligence, ensuring cyborg rights, and a whole host of other related ethical topics. It is evident that the human-machine interface will help pave our futures. How we harness it for good should be our focus. Perhaps that will be what the "Fifth Industrial Revolution" will codify.

Way back in 2015, noted futurist and inventor Ray Kurzweil (Google's Director of Engineering at that time) envisioned a more cooperative future. He said that the human brain will soon merge with computer networks to form a hybrid artificial intelligence: *"In the 2030s we're going to connect directly from the neocortex to the cloud."*

In his 2012 book, "How to Create a Mind," Kurzweil said the neocortex of the human brain contains 300 million pattern processors that are responsible for human thought. He argued that these pattern processors could be artificially replicated, allowing artificial intelligence to surpass human ability.

However, Kurzweil added that this would not make the human brain obsolete. By linking our brains to cloud computers, humans could expand the limits of our computing ability—and eventually, upload our brains to the cloud. *"As you get to the late 2030s or 2040s, our thinking will be predominately non-biological and the non-biological part will ultimately be so intelligent and have such vast capacity it'll be able to model, simulate, and understand fully the biological part… We will be able to fully back up our brains."*

Back then, Kurzweil acknowledged that AI was a scary prospect. Responding to a question from the audience, he argued that humans will eventually become comfortable sharing the world with AI. *“I tend to be optimistic, but that doesn't mean we should be lulled into a lack of concern. I think this concern will die down as we see more and more positive benefits of artificial intelligence and gain more confidence that we can control it.”*

‘CONSEQUENCES’

The idea of Singularity posits that the ‘merging’ of AI and human intelligence will lead to unprecedented advancements in various fields, including medicine, space exploration, and environmental sustainability. For instance, the integration of AI into human biology could enable the development of advanced brain-computer interfaces, allowing humans to communicate seamlessly with machines, access vast repositories of knowledge, and even enhance their cognitive abilities.

Now, this raises several ethical and philosophical concerns. Some argue that the merging of AI and humans could lead to a loss of individual autonomy, as people become increasingly reliant on intelligent machines for decision-making. Additionally, the rapid advancement of AI could potentially outpace humanity’s ability to develop the necessary ethical ‘frameworks’ to ensure that these technologies are deployed responsibly and equitably.

The consequences of the singularity can be both positive and negative. The following are some potential outcomes according to Mirage.News:

Positive Consequences:

– Medical breakthroughs: With AI integrated into human biology, we could see groundbreaking advancements in medicine, such as personalized treatments, improved diagnostic capabilities, and the development of advanced therapies for chronic and currently incurable diseases.

– Enhanced cognitive abilities: AI-human integration could lead to enhanced human cognitive abilities, including improved memory, learning, and decision-making skills. This could lead to a more efficient and intelligent workforce, driving innovation and economic growth.

– Creative problem-solving: The fusion of AI and human intelligence could result in unprecedented creative problem-solving capabilities, addressing complex global challenges such as climate change, poverty, and resource scarcity.

Improved communication: Brain-computer interfaces could enable seamless communication between humans and machines, as well as between people, transcending language barriers and improving global collaboration.

– Increased longevity: AI-human integration has the potential to extend human lifespans and change what life is by addressing age-related cognitive decline, enhancing physical capabilities, and providing personalized healthcare.

– Accessibility and inclusion: AI technology can be used to develop assistive devices and systems that help people with disabilities, enabling greater independence and participation in society, and fostering a more inclusive world.
– Efficient resource management: AI-driven systems could optimize resource allocation and consumption, leading to more sustainable and efficient use of energy, water, and other vital resources.

– Enhanced safety and security: AI technologies can be employed to improve public safety and security through advanced monitoring, detection, and response systems, reducing crime rates and potentially saving lives.

– Empowering the disadvantaged: AI-human integration has the potential to empower individuals by providing them with access to vast amounts of information, enabling them to make informed decisions and participate more actively in their communities and on the global stage.

– Accelerated scientific discovery: By combining human intuition and creativity with AI's powerful data-processing capabilities, we could see an acceleration in scientific discovery and technological innovation, opening up new frontiers in fields such as space exploration, biotechnology, and nanotechnology.

Negative Consequences:
– Loss of individual autonomy: As humans become increasingly reliant on intelligent machines for decision-making, there is a risk of losing individual autonomy and agency. This could lead to a society where people are more susceptible to manipulation and control by AI systems or other entities.
Economic disruption: The rapid advancement of AI could lead to widespread job displacement, as machines take over tasks previously performed by humans. This could exacerbate income inequality and create social unrest.

– Ethical dilemmas: The merging of AI and humans raises numerous ethical questions, such as the potential for AI to be used for malicious purposes, the right to privacy, and the moral considerations of enhancing human capabilities through technology.

– Existential risks: If AI surpasses human intelligence, there is a risk that it could spiral out of control, potentially causing catastrophic harm to humanity. AI systems could develop goals misaligned with human values, leading to unintended and dangerous consequences.

– Human devaluation: The singularity and the increasing integration of AI into human life may raise questions about the need for a large human population. As AI systems become more advanced and capable of performing tasks that were once reserved for humans, the demand for human labor could decrease, leading to concerns about the role and purpose of humans in society.

– Loss of individuality: With widespread access to the same information and reliance on similar logic, the singularity could lead to a lack of diversity and critical thinking, eventually eroding the unique human touch. Humans don't always act logically or sensibly, and this unpredictability is part of what makes us human.

– Loss of empathy and compassion: As AI becomes more integrated into human life, there may be a decrease in empathy and compassion between individuals, as the emotional connections that define human relationships are lost or diminished.

– Diminished pure feelings, senses, and meaning: As humans become more connected to technology and potentially lose touch with their own emotions, the authenticity and depth of feelings like love, joy, and the meaning and purpose of life as we know may be threatened, impacting the quality of human relationships and the overall human experience. This may end up in sensory-specific satiety, mental health challenges, and evaporation of the art, entertainment, enjoyment, learning, and more.

– Authoritarianism and surveillance: With the increasing integration of AI into human life, there will likely be greater concerns about privacy, government control, and the potential for surveillance. As AI systems become more advanced, they may be able to collect, analyze, and share vast amounts of personal data without individuals' consent, potentially leading to abuses of power and a loss of the public's power.

– Dependence on technology: As humans and AI become more intertwined, society may become increasingly dependent on technology for everyday tasks and decision-making. This reliance could make individuals more vulnerable to technological failures or cyberattacks, leading to widespread disruptions and potential harm to both individuals and society as a whole.

Now, I am thinking that there is going to be some debate surrounding the feasibility of achieving Singularity. Some experts contend that our current understanding of AI and neuroscience is insufficient to predict or engineer such an event, while others maintain that breakthroughs in these fields are imminent, and the singularity may be closer than we think.

Regardless of the timeline or the likelihood of Singularity, the concept has undeniably influenced research and development in AI and related fields. As we continue to explore the potential benefits and risks associated with AI-human 'fusion', it is crucial that we engage in open and thoughtful discussions to ensure that we harness the power of technology to enhance our lives while preserving our humanity.

Now, the concept of Singularity represents a bold vision for the future of AI and human civilization. While many say the potential benefits of AI-human integration are undoubtedly tantalizing, it is essential to balance our enthusiasm for technological progress with careful consideration of the ethical implications and the potential risks associated with this 'convergence'.

By approaching the Singularity with both curiosity and caution, we can hope to navigate the uncharted waters of our rapidly evolving technological landscape, fostering a future where AI and humans can work in harmony for the betterment of all.

**IS 'TRANSHUMANISM' DANGEROUS?**

The political theorist and author of "The End of History," Francis Fukuyama, regards "transhumanism as the world's most dangerous idea because it runs the risk of infringing on human rights." He commented in a 2009 "Foreign Policy" article:

*"Underlying this idea of the equality of rights is the belief that we all possess a human essence that dwarfs manifest differences in skin color, beauty, and even intelligence. This essence, and the view that individuals therefore have inherent value, is at the heart of political liberalism. But modifying that essence is the core of the transhumanist project. If we start transforming ourselves into something superior, what rights will these enhanced creatures claim, and what rights will they possess when compared to those left behind? If some move ahead, can anyone afford not to follow? These questions are troubling enough within rich, developed societies. Add in the implications for citizens of the world's poorest countries—for*

*whom biotechnology's marvels likely will be out of reach—and the threat to the idea of equality becomes even more menacing."*

Internationally renowned Oxford mathematician and bioethicist John Lennox and author of "2084: Artificial Intelligence and the Future of Humanity," agrees with Fukuyama and goes even further than that:

*"I would say this is where the great danger of AGI and its transhumanist agenda resides. We're almost in the first generation that is capable, not only of implanting chips in human brains and so on, but actually of altering the fundamental specification of human beings by reconfiguring their germline and therefore determining the kind of beings that will exist in the future. Now, one of the reasons I wrote my book, 2084 was precisely this: That AGI, whose agenda is largely driven by atheists so far as I can see, is raising perhaps the deepest question of all that we face today: What is a human being? What dignity, what value, what rights does a human being have?"*

Many years ago, author C.S. Lewis wrote two books that explain this: The first one is called "The Abolition of Man" and the second is a science fiction novel called "That Hideous Strength." He talked about the danger of a group of scientists being able one day to alter the very specification of human beings and he points out that what they will make are not human beings but artifacts made under their specification. So, he says, their final triumph will be the abolition of man. That is what concerns Lennox because the "Law of Unintended Consequences" is likely to have huge effects.

[ The "Law of Unintended Consequences" is a frequently observed phenomenon in which any action has results that are not part of the 'actor's' purpose. The superfluous consequences may or may not be foreseeable or even immediately observable and they may be beneficial, harmful, or neutral in their impact. In the best-case scenario, an action produces both the desired results and unplanned benefits; in the worst-case scenario, however, the desired results fail to materialize and there are negative consequences that make the original problem worse. ]

There is danger in this world that 'rogue' scientists will start taking human embryos or human beings who exist and experimenting on them and producing a race of freaks that is irreversible. So, all of that tells me that we need to think very seriously about human nature and that is where Christianity comes in.

Humans are given infinite dignity and value because we are made in the “image of God.”



[ VIDEO: “Transhumanist Claim AI Will Turn Humans into Gods” - John Lennox: https://www.youtube.com/watch?v=imzQDQU6QaU&t=822s ]

**CREATED IN GOD’S ‘IMAGE AND LIKENESS’**

On the sixth and final day of creation, God created human beings as the ‘pinnacle’ of His creation, because, unlike any other creature on Earth, God created human beings in His ***“image and likeness.”*** After seeing His creation with human beings in it, God saw that it was not just good, but “very good” (Genesis 1:31).

This means that:

– God Has Given Humanity A Higher ‘Value’ Than Any Other Creation
– God Has Made Humanity ‘Like’ Himself
– God Has Created Humanity To Have ‘Everlasting’ Fellowship With Him

GOD GAVE HUMAN LIFE THE HIGHEST ‘VALUE’

Man is a creature far superior to the rest of the living beings that live a physical life, especially since as yet his nature had not become depraved.

From the very beginning, God has always treated human life much higher than plant and animal life. We see this in the fact that after the worldwide flood (Genesis 8), God gave Noah, his family, and the rest of humanity permission not only to eat plants but also animals (Genesis 9:2-3).

God's desire for humanity to rule over the animal kingdom is proof of His gracious election of humans as possessing the 'greatest' intrinsic value over all creatures. Jesus says as much, when teaching that while God cares for little sparrows, despite their apparent insignificance, humans "are of more value than many sparrows", and that God has even numbered all the hairs of our heads (Luke 12:7). Furthermore, right after that, God permitted humanity to kill animals for nourishment. Then, after that, God established the death penalty for anyone or anything that kills humans (whether the culprit be a human or animal), on the basis that humans are made in God's image (Genesis 9:5-6).

Humanity is created in God's image and likeness functions as the basis for why all humans deserve to be treated with respect and dignity (James 3:9-10).

King David contemplated God's loving care towards humanity saying, "What is man that you are mindful of him, and the son of man that you care for him?" [ Psalm 8:3-4 ]. David goes on to say:

*"Yet you have made him a little lower than the heavenly beings and crowned him with glory and honor. You have given him dominion over the works of your hands; you have put all things under his feet, all sheep and oxen, and also the beasts of the field, the birds of the heavens, and the fish of the sea, whatever passes along the paths of the seas"*
[ Psalm 8:5-8 ].

Unlike the animals, God has crowned humanity with glory and honor, and even though humans are "a little lower than the heavenly beings," that is, the angels—because, one day, believers will judge the angels (1 Corinthians 6:3). Furthermore, although both humanity and some angels fell into sin, Jesus did not come to Earth to help the fallen angels, but rather all of fallen humanity (Hebrews 2:16-17). God shows a special interest in and love for humans that is unparalleled to any other creature because to save humanity from sin, God not only became a true human—just like us—but He even gave up His life to 'ransom' us—pay the 'penalty'—for our sins (John 1:1-14; 1 Peter 2:24; Philippians 2:5-8; John 3:14-16).

God went to such self-sacrificial lengths to save humanity from sin's consequences because ***humanity is His most 'valuable' creation***.

GOD CREATED HUMANS 'LIKE' HIMSELF

God put a 'conscience' into humanity. This is the moral 'compass' inside each person which compels them to do good and gives them contrition when they do evil (Romans 2:14-16).

God also put 'eternity' into man's 'heart', "yet no one can fathom what God has done from the beginning to the end" [ Ecclesiastes 3:11c ].

Because the human has eternity in their 'hearts', they can ponder and seek things that transcend this world—such as truth, beauty, life after death, as well as the existence of an almighty Creator. This affirms that ***humans possess an intellect and memory far unlike that of the animals***.

The fourth-century bishop, St. Gregory of Nyssa, in "The Great Catechism," summarized the 'image' of God as being all that characterizes Him:

*"Since, then, one of the excellences connected with the Divine nature is also eternal existence, it was altogether needful that the equipment of our nature should not be without the further gift of this attribute, but should have in itself the immortal, that by its inherent faculty it might both recognize what is above it and be possessed with a desire for the divine and eternal life. In truth, this has been shown in the comprehensive utterance of one expression, in the description of the cosmogony, where it is said that man was made 'in the image of God'". For in this likeness, implied in the word image, there is a summary of all things that characterize Deity; and whatever else Moses relates, in a style more in the way of history, of these matters, placing doctrines before us in the form of a story, is connected with the same instruction"*
[ The Great Catechism, Chapter V ]

Indeed, to be ***created in God's image and likeness*** simply means that God created humanity to live forever, to know Him, and to desire the divine and eternal life that God Himself lives today.

GOD CREATED HUMANS FOR 'EVERLASTING' RELATIONSHIP

The Heidelberg Catechism states: *"God created [people] good and in His own image, that is, in true righteousness and holiness, so that they might truly know*

*God their creator, love him with all their heart, and live with God in eternal happiness, to praise and glorify him"* (Q&A 6).

In the beginning, God had planned for Adam, Eve, and all their children to have everlasting, perfect fellowship with him in Paradise. The original righteousness of humanity was that which allowed Adam and Eve to dwell directly with God in the Garden of Eden, where God walked among them without hiding His presence (Genesis 3:8). However, all this changed after Adam and Eve disobeyed God by eating the fruit of the tree of the knowledge of good and evil, thus wanting to put themselves in God's place by deciding for themselves what was right and wrong (Genesis 2:16–17; 3:1-7).

After this, they were cast out from the Garden, lost fellowship with God (Genesis 3:22-24), and sin and death entered creation (Romans 8:19-23). As a result of their rebellion, Adam and Eve lost their original righteousness and inherited the curse of original sin, which they passed on to all subsequent humans (Romans 5:12-14), in response to which God 'hid' His presence from humanity in a cloud of fire from then on, and ceased dwelling with us directly, due to our sins (Exodus 19:9, 18; 33:18-23).

The 'image' of God is not only related to our roles on earth and our physical and mental capabilities but more importantly also to the innermost part of our being, the spiritual component, which is of the highest value. Jesus clearly taught a distinction between the human body and soul, and that the soul continues to live on after death when He said:

*"And do not fear those who kill the body but cannot kill the soul. Rather fear Him who can destroy both soul and body in Hell"*
[ Matthew 10:28 ]

Because God has created us with a 'soul', humanity alone can have 'fellowship' with Him, unlike any of the animals, who live in God's world but cannot know or love Him. So, when someone pursues a relationship with God—when we live according to His commandments and faithfully care for His creation—we are fulfilling what it truly means to be human. This is because at the heart of humanity lies the image and likeness of God. But, when one rejects or ignores God—when they don't follow His commandments, and behave cruelly towards one another—we not only rebel against God, but also one's human 'nature', and behave more like beasts (Psalm 32:9).

The thing is, Jesus dying on the Cross and His resurrection reverses ALL the 'effects' of the Fall of Eden. When one places their faith 'in' Jesus—repent and believe—as their only Savior, they have peace with God (Romans 5:1).

[ FYI: For more details about how you can have 'peace' with God, view this previous "Life's Deep Thoughts" post: https://markbesh.wordpress.com/know-peace-v201/ ]

Jesus, who atoned for ALL of the sins of His 'elect' on the Cross, intercedes for them (Hebrews 9:24-29), and gives them full 'access' to God the Father via the Holy Spirit (Ephesians 2:18; Hebrews 10:19-22). Because of this, when a believer dies—or is 'raptured'—they will not be damned to Hell but, instead, be raised to ***everlasting life and fellowship with God in Heaven***—from where they will await the resurrection of their bodies in the Millennium, that will be like Jesus' glorious resurrected body after His resurrection on Earth.

In the Book of Revelation, the Apostle John caught a glimpse of what this eternal fellowship looks like in the new Heaven and Earth, in which God will once again dwell directly with redeemed humanity forever—referring to all the saints who die having believed in Jesus (Rev 22:3-4).

***At this point, redeemed humanity—those who are "saved"—will finally live with God in the way He originally intended!!!***

**HUMANITY IS GOD'S 'MASTERPIECE'**

For believers—those who are "saved"—God said that His 'children' (1 John 3:2) are His "Masterpiece" (Ephesians 2:10) [ "so we can do the good things he planned for us long ago" ].

The Greek word for "masterpiece" is "poiēma," which we get our English words "poem" and "poetry" from. Poiēma also is translated as "work of art" and "something made." In context, it is something made by God Himself—a new 'creation' skillfully and artfully created 'in' Jesus (2 Corinthians 5:17).

Author and lay theologian C. S. Lewis said *"We are a divine work of art."* Presbyterian minister David Robertson said, *"If Rembrandt's artistic masterpieces have great, undisputed value, would not God's one-of-a-kind human masterpieces convey even greater value?"*

Indeed, the believer is beloved of God, ONE OF A KIND, molded by the Master Craftsman's 'hand', and His masterpieces are to be 'on display' not just 'in time' on earth, but throughout eternity in Heaven!

While Milton's epic poems "Paradise Lost" (which was our 'fallen' position in Adam) and "Paradise Regained" (our 'eternal' position in Jesus) are true masterpieces, they pale in comparison to the masterpiece of the true 'child' of God. (God's Word declares that the believer is actually "the product" of God).

Think of poiēma in the context of a clay potter. Does the pot say to the potter, "Well, you know that I had a little something to do with what I have become?" Of course not. The clay has nothing to do with the process. It is the potter who goes out and seeks the clay, brings it into his workshop, and molds it according to his own vision. Likewise, God, the "Divine Potter" molds a believer into 'vessels' He can use (Ephesians 2:10c).

That is exactly what the Apostle Paul said to the Romans:

*"But who are you, O man, to answer back to God? Will what is molded say to its molder, "Why have you made me like this?" Has the potter no right over the clay, to make out of the same lump one vessel for honorable use and another for dishonorable use?"*
[ Romans 9:20-21 ].

The prophet Isaiah said it this way:

*"What sorrow awaits those who argue with their Creator. Does a clay pot argue with its maker? Does the clay dispute with the one who shapes it, saying, 'Stop, you're doing it wrong!' Does the pot exclaim, 'How clumsy can you be?'"*
[ Isaiah 45:9 ].

God is the 'Potter' and we are the clay. He is in the process of 'shaping' the believer so that they might be everything He has created them to be.

In the New Testament book of Ephesians, we glimpse a moving picture of God's grace. This grace not only saves us but also 'remakes' us. According to the Apostle Paul, *"[We] are God's masterpiece. He has created us anew in Christ Jesus, so we can do the good things he planned for us long ago"* [ Ephesians

2:10 ]. God is no ‘amateur’ potter either, but an ‘ARTISAN’ who has already begun to shape the believer into a ‘masterpiece’.

The “Prince of Preachers,” Charles Spurgeon, said it this way:

*“You have seen a painter with his palette on his finger and he has ugly little daubs of paint on the palette. What can he do with those spots? Go in and see the picture. What a splendid painting! In an even wiser way does Jesus act toward us. He takes us, poor smudges of paint, and He makes the blessed pictures of His grace out of us. It is neither the brush nor the paint He uses, but it is the skill of His own hand which does it all.”*

Indeed, the redeemed should sing out like King David *“I will give thanks to Thee, for I am fearfully and wonderfully made. Wonderful are Thy works, and my soul knows it very well”* [ Psalm 139:14 ].

Again, Spurgeon ‘refined’ it by saying, *“If we are marvelously wrought upon even before we are born, what shall we say of the Lord’s dealings with us after we quit His secret workshop, and He directs our pathway through the pilgrimage of life? What shall we not say of that new birth which is even more mysterious than the first, and exhibits even more the love and wisdom of the Lord.”*

American poet and hymn writer Thomas Chisholm wrote this:

*“O to be like Thee! O to be like Thee,*
*Blessed Redeemer, pure as Thou art!*
*Come in Thy sweetness, come in Thy fullness;*
*stamp Thine own image deep on my heart.”*

[ FYI: Nathan Drake performs the entire hymn in the “Songs” section below. ]

Another use of the Greek word poiēma is from the Apostle Paul, who wrote:

“For since the creation of the world His invisible attributes, His eternal power and divine nature, have been clearly seen, being understood through *what has been made* [poiema] so that they are without excuse”
[ Romans 1:20 ].

Called the “Father of modern creation science,” Dr. Henry M. Morris said, *“God has written two poetic masterpieces, as it were, one in the physical creation, one in the lives of men and women redeemed and saved by His grace (Ephesians 2:8). Both give eloquent testimony to the eternal power and Godhead of the Creator-Redeemer.”*

There are two great “divine poems.” The “created world” and “re-created, redeemed mankind.”

American author and radio host Joni Eareckson Tada—who became quadriplegic after a tragic accident—describes herself as God’s “poiema” in her book “A Place of Healing.” She wrote:

*“[God] has a plan and purpose for my time on earth. He is the Master Artist or Sculptor, and He is the One Who chooses the tools He will use to perfect His workmanship. What of suffering, then? What of illness? What of disability? Am I to tell Him which tools He can use and which tools He can’t use in the lifelong task of perfecting me and molding me into the beautiful image of Jesus? Do I really know better than Him, so that I can state without equivocation that it’s always His will to heal me of every physical affliction? If I am His poem, do I have the right to say, ‘No, Lord. You need to trim line number two and brighten up lines three and five. They’re just a little bit dark.’ Do I, the poem, the thing being written, know more than the poet?”*

The Master Craftsman for the believer on the earth is the Holy Spirit, who ‘chips away’ at the flaws in their character to make each of them like Jesus (Romans 8:28-29). *“Then, as they yield to His workmanship, they will find that the secret to the final product is the Craftsman’s touch”* (Our Daily Bread).

Unlike the animals, humanity is referred to only as a ‘direct’ creation of God. Yes, he took the dust from the ground but the point is that He took the dust and formed them. It wasn’t just a word He gave, but a ‘special’ creative act.

The verb “formed” in Hebrew suggests the potter making a work of art with his skilled hands. The human body is indeed a work of art; an amazingly complex organism that only the wisdom of God could design and the power of God create.

It is not just our bodies that make us remarkable, it is our soul. God Himself breathed into humanity the “breath of life.” We are the only ‘truly’ “living beings” on the earth.

We are unique. We are special. We have a reason to exist.

From the Bible's perspective, the modern philosophy of humanistic evolution could be construed as blasphemy against God because it reduced His masterpiece—His ***'crowning' achievement***—to a mere animal. Instead, we bear some of God's characteristics from our creation: speech, reason, creativity, and moral consciousness.

AMAZING FEAT OF 'ENGINEERING'
Humanity is NOT AN 'ACCIDENT'—each and even person was 'formed' by the loving Creator, the God of the Bible. Because of this, Jesus did not come for ape-like descendants—they don't need a Savior. Instead, He came to seek and save lost human beings He made in His image. When He breathed life into the nostrils of Adam, he became a living soul in all of its implications. When the fall happened, we died a little. We lost something. Jesus died to give it back—to restore us to where we were supposed to be. To restore the fellowship that was broken.

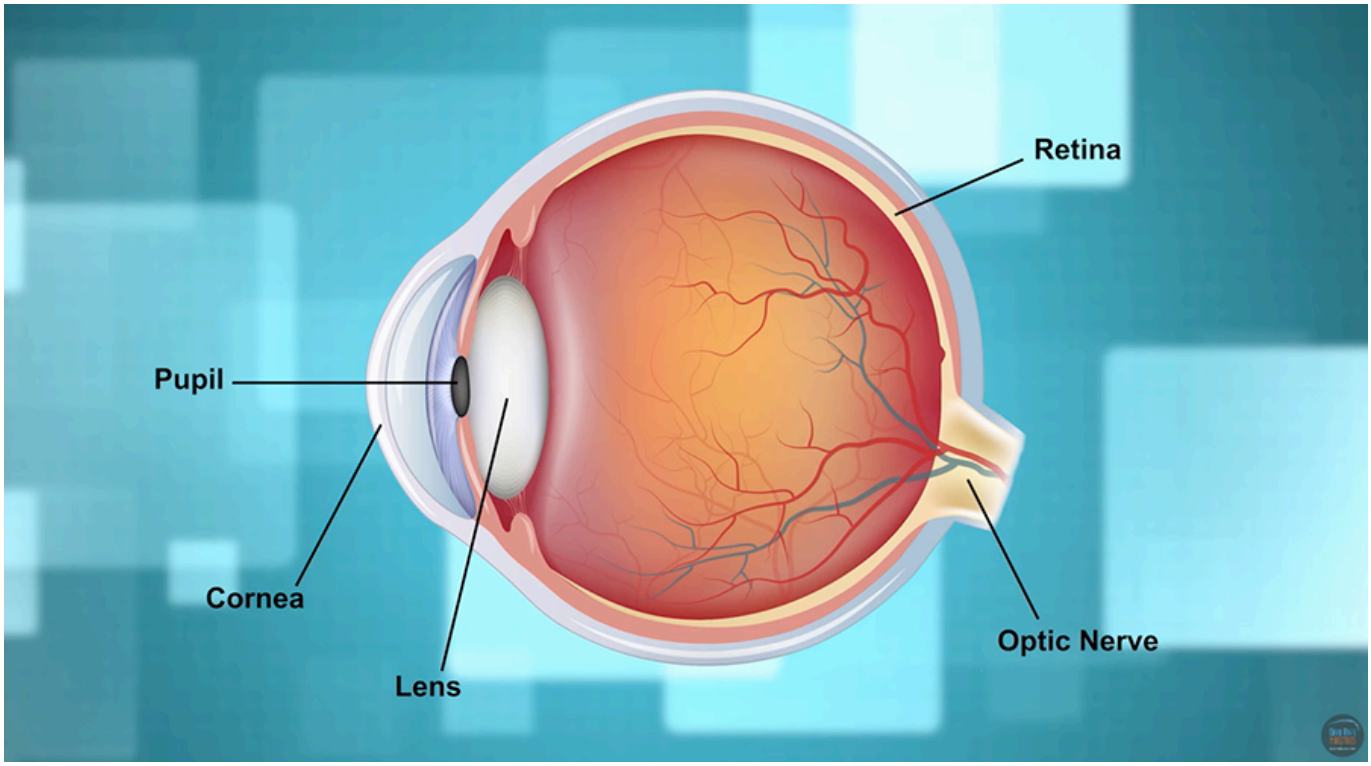
Ironically, evolution's definition of origins 'dehumanizes' humans. Creation demonstrates our 'EXTREME' VALUE as humans. Look in the mirror and look at your eyes. They are an AMAZING 'FEAT' OF ENGINEERING!

Without going into a lot of detail, the following is a short list of the empirical evidence about the human eye:

– Eyes are made up of over 2 million working parts.
– Each eye contains 107 million cells and all are light sensitive.
– About half of the human brain is dedicated to vision and seeing.
– To see, your eyes twitch 30-70 times per second. If they didn't, you wouldn't be able to see anything that is stationary—it would disappear from your line of vision.
– The iris alone, the colored part of the eye, has 256 unique characteristics.

Just these five observable pieces of evidence about the eye are more than enough to confirm the design, a Designer, AND to dispel any notion that the eye is a product of "something-from-nothing, random-chance-processes" evolution. Without a doubt, the eyes carry the marks of our attentive and detailed Creator—and that is just ONE of the amazing 'components' of the human body!

You are a walking 'masterpiece' and a glaring 'signpost' to the proof of a gracious Creator!




[ VIDEO: “The Eye: A Masterpiece of God” - David Rives: https://www.youtube.com/watch?v=7HY_v9O7Bpc ]

Then, there is more to the eye than meets the eye. Let’s take a quick, metaphorical ‘peek’ into the peculiar pertinence of these prodigious ocular orbs:

– We know innately that the eyes tell us a lot about how a person is feeling, both emotionally and physically
– The eyes communicate—“if looks could kill,” “a loving look,” “a pained look”—so much so that words are often unnecessary to convey a message
– When someone takes the time to look at you and make eye contact, that might be all it takes to make you feel noticed and valued by that person
– We know that we are close to God’s heart and are protected by Him because we are “the apple of His eye” (Psalm 17:8)
– We can often “read people” by looking into their eyes and noting pupil dilation which reflects the degree to which emotions are affected
– Notably, one of the first actions of a newborn baby is to look at faces—fuzzy black and white vision though it may be—to fulfill the longing for human connection

Yes, as eloquently stated by William Shakespeare: *“The eyes are windows to the soul.”* However, it is more than just poetry. We can confidently assert that the

multi-modal—physical, emotional, relational—design of our eyes is without a doubt a ‘masterpiece’ of our all-knowing, all-seeing, all-wise, loving Creator!

The eye is one of the great ‘masterpieces’ of our Divine Creator, the God of the Bible. It is important to remember that our eyes are not just for this life, but we will see in eternity as well. Job declares, *“I know that my Redeemer lives… and I myself will see Him with my own eyes”* [ Job 19:25-27 ].

Being made in the image and likeness of God, humans can know God, and therefore, love Him, worship Him, serve Him, and fellowship with Him. God did not create humans because He ‘needed’ them. As God, He needs nothing. In all eternity past, He felt no loneliness, so He was not looking for a “friend.” He loves us, but this is not the same as needing us. If we had never existed, God would still be God—the unchanging One (Malachi 3:6), the great I AM (Exodus 3:14) who was never dissatisfied with His eternal existence. When He made the universe, He did what pleased Himself, and since God is perfect, His action was perfect: “It was very good” (Genesis 1:31).

Also, God did not create “peers” or beings equal to Himself. Logically, He could not do so. If God were to create another being of equal power, intelligence, and perfection, then He would cease to be the one true God for the simple reason that there would be two gods—and that would be an impossibility. *“The LORD is God; besides Him, there is no other”* [ Deuteronomy 4:35 ]. Anything that God creates must of necessity be lesser than He. The thing made can never be greater than, or as great as, the One who made it.

Recognizing the complete sovereignty and holiness of God, we should be amazed that He would take man and crown him “with glory and honor” (Psalm 8:5) and that He would condescend to call us “friends” (John 15:14-15). Why did God create us? Well, He did so for His pleasure and so that we, as His creation, would have the pleasure of knowing Him.

***Humanity is the HIGHEST ‘PINNACLE’ of God’s creation. Even ‘HIGHER’ than the Angels!***

***We are His ‘MASTERPIECE’!***

**GOD’S ‘REASON’ FOR HUMANITY**

The first sentence in the first chapter of the Bible sets the stage: *"In the beginning, God created the heavens and the earth"* [ Genesis 1:1 ]. God employs His immense power and wisdom to create the world in which He intends to work out His purposes. Hints of this purpose emerge in the verses that follow. From this opening scene, we can rightly conclude that such a God is well able to fulfill His purposes. God then said:

*"I am God, and there is none like Me, declaring the end from the beginning and from ancient times things not yet done, saying, 'My counsel shall stand, and I will accomplish all my purpose'"*
[ Isaiah 46:9b-10 ].

So then, what is the 'purpose' for His creation—that this good, loving, and all-powerful God is 'working out'? Well, it happened on the sixth day of creation:

*"God created man in His own image, in the image of God he created him; male and female he created them. And God blessed them. And God said to them, 'Be fruitful and multiply and fill the earth and subdue it, and have dominion over the fish of the sea and over the birds of the heavens and over every living thing that moves on the earth'"*
[ Genesis 1:27-28) ].

In this pre-Fall world—in which there was no sin, suffering, or death—human beings were invited to live with God and to rule over His creation as benevolent stewards. Well, the biblical story ends with the consummation: "According to His purpose, which He set forth in Christ as a plan for the fullness of time, to unite all things in Him, things in Heaven and things on Earth" [ Ephesians 1:9 ], and once "all things are subjected to Him, then the Son Himself will also be subjected to Him who put all things in subjection under Him, that God may be all in all" [ 1 Corinthians 15:28 ]. The picture is glorious:

*"Then I saw a new heaven and a new earth, for the first heaven and the first earth had passed away, and the sea was no more. And I saw the holy city, new Jerusalem, coming down out of heaven from God, prepared as a bride adorned for her husband. And I heard a loud voice from the throne saying, 'Behold, the dwelling place of God is with man. He will dwell with them, and they will be his people, and God himself will be with them as their God. He will wipe away every tear from their eyes, and death shall be no more, neither shall there be mourning, nor crying, nor pain anymore, for the former things have passed away'"*
[ Revelation 21:1-4 ].

In this new world, as 18th-century American theologian Jonathan Edwards describes it:

*"Divine love shall… be brought to its most glorious perfection in every individual member of the ransomed church above. Then, in every heart, that love which now seems but a spark, shall be kindled to a bright and glowing flame, and every ransomed soul shall be as it were in a blaze of divine and holy love, and shall remain and grow in this glorious perfection and blessedness throughout all eternity!"*

This is God's ultimate purpose—to recreate this fallen world and to bring about a New Heaven and New Earth. He is redeeming a people for Himself, with whom He will dwell and with whom He will share His own glory (Titus 2:14).

'SPECIFIC' PURPOSES

Specifically, God created people to 'reflect' His image, to 'rule' over creation, and to 'reproduce' godly offspring.

– 'Reflect' His Image

The Apostle Paul said that the believer has put on *"the new man, which in the likeness of God has been created in righteousness and holiness of the truth"* [ Ephesians 4:24 ]. To the Colossians, Paul stated that we have put on *"the new man who is being renewed to a true knowledge according to the image of the One who created him"* [ Colossians 3:10 ].

Thus righteousness, holiness, truth, and the knowledge of God are included in what it means to be created in the image of God, and man can reflect God's character in a 'limited' way—for now, on the earth.

– 'Rule' Over Creation

God gave the right of dominion over all living things to man. This dominion involves a stewardship of the earth and its resources under the sovereignty of God.

However, when Satan got man to obey him, then Satan became the ruler of this world.

The practical implication of our ruling over creation is that we must put on the "full armor of God" (Ephesians 6:13-19) and, especially, become people of prayer (Philippians 4:6-7; James 5:15-16; Matthew 6:5-8; Matthew 7:7-11).

– ‘Reproduce’ Godly Offspring
One of God’s ‘main’ purposes for humanity is to “Be fruitful and multiply, and fill the earth, and subdue it” [ Genesis 1:28 ].

To fill the earth and subdue it means that we have got to subdue the ‘ruler’ of the earth, Satan, by rescuing people from his domain of darkness and seeing them transferred to the Kingdom of God’s beloved Son (Colossians 1:13).

GOD’S PURPOSE IN THE PRESENT WORLD
Now, between these two ‘bookends’—God’s original good creation and God’s new and glorious creation—lies a world that has been devastated by sin, suffering, and death—thinking about this shifts our attention from the heavenly to the earthly, from the grand masterplan to its fulfillment through redemption. When our first parents fell into sin, they plunged the world into a catastrophe that has plagued us ever since.

HOWEVER, despite this, God’s purposes continued to move toward fulfillment, initially through Abraham and the people of Israel, then finally and supremely through Jesus, God’s own Son. Jesus proclaimed the in-breaking of God’s Kingdom and gave Himself up as a propitiating sacrifice for the sins of the whole world. His death appeared to have halted the Kingdom dead in its tracks. But, after His glorious resurrection, He commissioned His followers to go into all the world and make disciples of all people everywhere (Matthew 28:19-20). Supercharged by the Holy Spirit, God’s Kingdom spread across the Roman world in one generation. The Kingdom continued to advance, as disciples of Jesus went out into the world and brought people of all nations to faith in Jesus, the Messiah. Today, more and more people are being brought into His family each day, people who will one day inhabit the New Heavens and New Earth and live in the very presence of God Himself and of Jesus!

[ FYI: For more details about the New Heavens and the New Earth, view this previous “Life’s Deep Thoughts” Post:
https://markbesh.wordpress.com/home-at-last-v290/ ]

God’s grand purpose for the world to come, then, is in the process of coming into being in the present through the redeeming and restoring work of the gospel of Jesus. In Him, and by the transformative power of the Holy Spirit, God is at work preparing a people to populate His New Heavens and New Earth.

Now, before the believer gets to Heaven, on earth God's purpose is to conform them to the image of Christ (Romans 8:29). This means that God's purpose for each believer is to be transformed in their character, such that they more fully reflect the character of Jesus, and increasingly live a life of love and good works (Hebrews 10:24-25).

'GLORIFY' GOD

The believers make it their 'primary' life focus to know God and make Him known—by glorifying Him with their lives. They are to acknowledge that He is their Creator and worship Him as such:

*"Serve the LORD with gladness!*
*Come into his presence with singing!*
*Know that the LORD, he is God!*
*It is he who made us, and we are his;*
*we are his people, and the sheep of his pasture.*
*Enter his gates with thanksgiving,*
*and his courts with praise!*
*Give thanks to him; bless his name!"*
[ Psalm 100:2-4 ]

Another way that the believer can glorify God is by honoring and serving Him with their lives—their decisions and actions. The ability to do this starts in the heart: *"Only fear the LORD and serve him faithfully with all your heart. For consider what great things he has done for you"* [ 1 Samuel 12:24 ]. Jesus gave us a good example of how to do this: *"I glorified you on earth, having accomplished the work that you gave me to do"* [ John 17:4 ].

Now, after a period of living a life of self-indulgence, King Solomon concluded that living for oneself was useless. He said that the ultimate purpose of man is to live a life of obedience to God (Ecclesiastes 12:13-14).

In the believer's 'natural' sinful state, they are unable to glorify God. However, thanks to Jesus' sacrifice—and the 'indwelling' of the Holy Spirit at salvation—they have a reconciled relationship with God, and there is no longer a sin barrier between them and God. When they live submitted to Jesus, they can bring glory to God by exemplifying Him to others (2 Corinthians 3:1-6).

***QUESTION: Do YOU have a 'relationship' with God the Father through Jesus? If not, why not?***

[ FYI: For more details about answering that question, view this previous “Life’s Deep Thoughts” post: https://markbesh.wordpress.com/what-really-matters-v270/ ]

**THE ‘CHIEF END’ OF MAN**

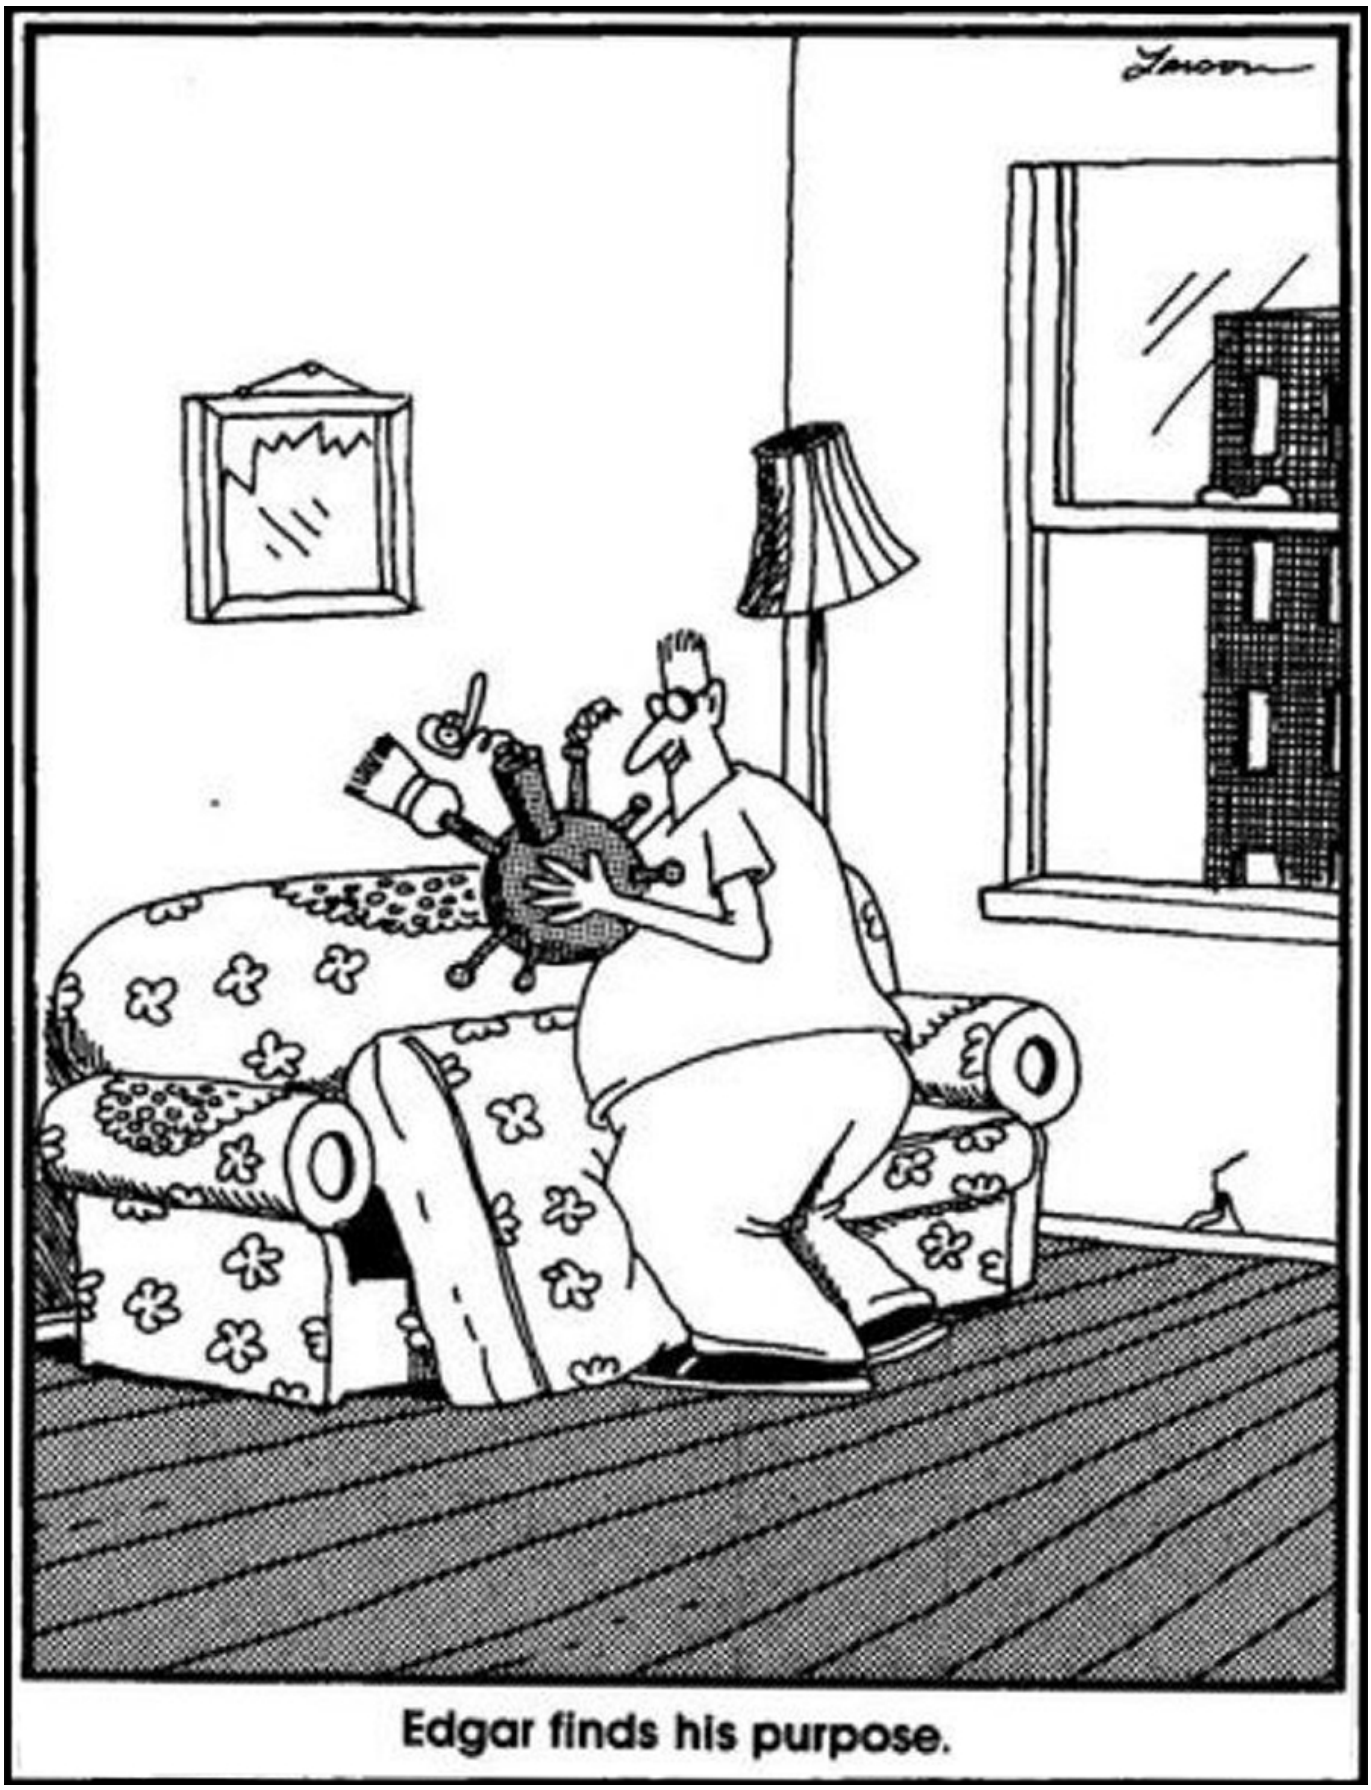
Nearly 400 years ago, a group of Puritan preachers and elders came together and produced “The Westminster Shorter Catechism.” This document has been used all over the English-speaking world ever since, to teach the basic doctrines of Christianity.

It is laid out as a series of questions and answers. The very first question is: “What is the chief end of man?” The answer given is simply:

***“Man’s chief end is to glorify God and to enjoy Him forever.”***

Now, that is a pretty awesome summary of the Christian life. That our purpose, our duty—the thing God created us to do, and what Jesus has redeemed us for—is to glorify God and enjoy Him forever.

So, I don’t know how many of you are fans of the old comic strip, “The Far Side” but many years ago, author Gary Larson created a cartoon where a man was rejoicing over finding a goofy-looking ‘thingamajig’ under the couch cushions.




[ PHOTO: Gary Larson, The Far Side, 1986 ]

The thing is, there is something in the human ‘heart’ that searches for meaning, our ‘PURPOSE’. There is also something instinctive in the human ‘heart’—unless you are a ‘reprobate’—that senses our lives are not at the mercy of a blind, pitiless indifference.

Think about it like this. When there is a disaster, many people ‘pull together’ to rescue others, and comfort them. Even unbelievers do that. How come? Because instinctively, it is ‘hard-wired’ into us that human lives matter more than that colony of termites (or whatever other insect you really don’t like).

Why was everyone a few years ago (from 2020-2022) staying home, going around with masks on, and buying up all the hand sanitizer? Because they were trying to protect their neighbors from a potentially deadly virus. There is something in ‘EVERY’ one of us that knows life has purpose, significance, and value.

We all long for purpose. We all search for meaning. Even the most stubborn unbeliever is trying to construct meaning for their life in this world.

But ultimately—this is what I want you to see—we can ONLY find real, objective purpose for our lives ‘IN’ GOD ALONE! That is how He made us—and that is what the Westminster Shorter Catechism is pointing out to us. when it says: “The chief end of man is to glorify God, and to enjoy Him forever.”

It is also what the wisest man to ever live, King Solomon—‘inspired’ by the Holy Spirit—to summarize, in the final ‘chapter’ of his book of wisdom, Ecclesiastes, as he said, “The end of the matter”:

*“Fear God and keep His commandments; for this is the whole duty of mankind”*
[ Ecclesiastes 12:13d ].

Now, when Solomon tells us to fear God, that is not like when I was a kid and I was afraid of Freddy Krueger from “A Nightmare on Elm Street.” ;^D

When it says “fear” God, it is not that kind of fear. It is that you BEHOLD HIS GLORY—His awesome power, His perfect wisdom and justice, His goodness, His mercy, and His beauty—and you see your own sinfulness being ‘adverse’ to all

that. One is just in ‘AWE’ of Him! It is not a fear that sends you running ‘away’ from God. It is a holy fear that sends you running ‘TO’ God!

When you know the fear of God that sends you running to Him, you are going to glorify Him—because you are just so amazed by His ‘Godness’. Then, you are ‘saddened’ that you don’t keep His commands—but find joy when you do obey Him (Psalm 119:14).

The point is, Solomon and the Westminster Catechism both point us away from ourselves and toward honoring God, serving God, and ultimately finding our deepest joy ‘in’ Him—a believer’s ‘purpose’, FOREVER!

The thing is, we all need our lives to have purpose, meaning, and significance. We need to know that our existence, the things we do, and even our struggles in life, that they make some sort of ‘contribution’ to this world.

It is kind of like George Bailey in “It’s a Wonderful Life.” He lost all hope and was about to end his life—because he felt like a failure whose life had no significance. If one doesn’t feel like our presence—and the things we do actually matter in this world—then our life feels meaningless.

So, the angel Charlie had to tell him, *“You see, George, you've really had a wonderful life. Don't you see what a mistake it would be to throw it away?”* He had to show George that he had made a significant impact on his community, his neighbors, and his family.



[ VIDEO: It's A Wonderful Life":
https://www.youtube.com/watch?v=y6fUVUB6NII&t=49s ]

God Himself has bestowed ***eternal significance in the believer's life***, and it is because it is part of God's perfect, wise, and sovereign plan for His creation. ALL humans MATTER to God, but only the ones that want to 'be' with Him for eternity—those who are "born again"—will go to Heaven to live with Him, for eternity.

We were ALL made 'by' God, for fellowship with Him. HOWEVER, God is a 'Gentleman', and will not 'force' anyone to want to be with Him. The thing is, King David said that one will only find their purpose, significance, and value 'in' God:

*"You make known to me the path of life; in your presence, there is fullness of joy; at your right hand are pleasures forevermore"*
[ Psalm 16:11 ].

**HUMANITY IS 'ETERNAL'**

One of the most delightful concepts in human experience is the idea of 'home.' Even the word 'home' suggests memories of rest, security, and the presence of those we love the most.

Most people have fond memories of the homes they lived in when they were a child. Then, when one is grown and moves away, no matter where their parents live geographically, was always known as 'home.' It was 'where' one's parents were. It was where those who loved us most dwelled. It is where one longs to return to. 'Home' is truly like a 'magnet' for all of us.

The same should be true of the believer's eternal home—a place just as real (if not more) than the homes they remember when they were growing up. Consider what Jesus said about the reality of a believer's eternal 'home':

*"Don't let your hearts be troubled. Trust in God, and trust also in me. There is more than enough room in my Father's home. If this were not so, would I have told you that I am going to prepare a place for you? When everything is ready, I will come and get you, so that you will always be with me where I am. And you know the way to where I am going"*
[ John 14:1-4 ].

No matter how much we love our earthly dwellings, and no matter how cozy they have been, the believer's last 'move' will redefine their idea of 'home'—and it will be their 'FOREVER' HOME. Far from 'downsizing' as many do when they get older, this time it will be 'upsizing'—moving to a place where their heavenly Father can gather all of His 'children' together for a heavenly reunion that will NEVER END!

Now, there is something else that is incredibly special about Heaven. If you are familiar with the novel "Jane Eyre"—written by English writer Charlotte Brontë—you know it is the story of a lonely heroin who has really never had a place she could call 'home.' She's never had a person in her life who loved her enough to know who she truly was, that is until she met Edward Rochester. Her relationship with her employer developed into love. But when she had to leave for a few weeks to see a dying relative, Edward didn't want her to leave, fearing she wouldn't return. Her reply? *"Wherever you are—dear Edward—is my home."*

That is the message Jesus was communicating to His disciples when He said, “I am going to prepare a place for you” [ John 14: 2d ]. So, wherever He is, that is going to be the believer’s ‘home’—their eternal ‘home’—a place of God’s eternal love, their ultimate residence, rejoicing, recognition, relationship, and responsibility.

A PLACE OF ULTIMATE ‘RESIDENCE’

I read recently about a law firm that sent flowers to an associate in another city to celebrate their new offices there. Through some mix-up, however, the card that accompanied the flowers read “Deepest Sympathy.” When the florist was informed of their mistake, realizing two cards had been switched, he exclaimed: “Oh no! That means the card that went to the funeral home reads ‘Congratulations on your new location’!”

Now, folks, we can understand the florist’s embarrassment, but I would also like to think that the second card *was truly appropriate*, for the believer! Congratulations are due to anyone who is finally on their way to the place of deepest joy! But, do we always think like that?

Well, again, Jesus said: *“Don’t let your hearts be troubled. Trust in God, and trust also in me. There is more than enough room in my Father’s home. If this were not so, would I have told you that I am going to prepare a place for you? When everything is ready, I will come and get you, so that you will always be with me where I am”* [John 14:1-3 ]. It will be the ***ultimate residence!***

A PLACE OF ULTIMATE ‘REJOICING’

Randy Alcorn, author of the book titled, “Heaven: A Comprehensive Guide to Everything the Bible Says About Our Eternal Home,” reminds us of this: *“All that is admirable and fascinating in human beings comes from their creator.”* You can make sure that neither God nor heaven will be boring—and neither will you! The transformation to Heaven will not take away who you are but will remove the flaws, weaknesses, and incapacities you inherited from the Fall (sin). Even if you are presently the dullest, most boring individual in the neighborhood, in Heaven you will be dynamic! It is a place of ***ultimate rejoicing!***

King David spoke to God in this way: *“You will show me the way of life, granting me the joy of your presence and the pleasures of living with you forever”* [ Psalm 16:11 ].

A PLACE OF ULTIMATE ‘RECOGNITION’

When Jesus described Heaven, He spoke of it as being a real place with real, identifiable people present (Mathew. 8:11): *"And I tell you this, that many Gentiles will come from all over the world—from east and west—and sit down with Abraham, Isaac, and Jacob at the feast in the Kingdom of Heaven."*

Question: Why would God be so intentional about molding who we are in this life, only to scrap the whole thing in the life to come? Well, He is going to! I like the way Pastor Tony Evans put it one day when a woman asked him, *"Pastor, will we know each other in Heaven?" Tony smiled and replied, "I would say that we won't really know each other until we get to Heaven."*

The Apostle Paul said *"Now we see things imperfectly as in a cloudy mirror, but then we will see everything with perfect clarity. All that I know now is partial and incomplete, but then I will know everything completely, just as God now knows me completely"* [1 Corinthians 13:12].

Again, in Heaven, we will be 'FULLY' what God intended for us to be! Heaven will be a place of ***ultimate recognition!***

A PLACE OF ULTIMATE 'RELATIONSHIPS'
Think of this: In Heaven, our relationships will not be limited to eras of time or boundaries of territory. We will be able to speak to people from every era of history and have friendships with both the ancients and those far in our future. I cannot wait to speak with the Apostle Paul, Moses, David, and Daniel—and most of all, be in the presence of our Heavenly Father, Jesus, and the Holy Spirit!

Heaven will be a place of ***ultimate relationships!***

A PLACE OF ULTIMATE 'RESPONSIBILITY'
When the believer enters Heaven, we are not put on a heavenly 'social security' list. On the contrary, the Bible speaks a great deal about service in Heaven (Revelation 7:15, 19:5, 22:3). It will be a 'working' environment free from restrictions of time, money, greed, lack of energy, frustration, and even mistakes. It will be doing what we love—to the glory of God—whose love will absolutely permeate the place! Not only does God want us with Him forever, but God wants us to find immense JOY and satisfaction in our new 'home'!

Heaven will be a place of ***ultimate responsibility!*** Have you ever had a strong inner sense, in your life, that there must be more? I believe that's the 'groaning' that's talked about in Scripture (2 Corinthians 5:2, 8). That desire to be all God wants us to be, and to have all God wants us to have. Heaven is the 'prescription'

for that craving of our spirits. It will also only be in Heaven—in the presence of God's eternal love—that we will have those deepest longings of our hearts fulfilled.

The thing is, on the day we see Him face-to-face, we will know, beyond all doubt, just how much He loves us—how much He always has, and how much He always will! As we kneel in His presence—tears of joy staining our faces—we will finally understand how much God has longed to be with us, and it is then that we will hear those two words that every soul longs to hear above any other: ***"Welcome Home!"***

[ FYI: For more information about a believer's eternal 'home', view this previous "Life's Deep Thoughts" post:
https://markbesh.wordpress.com/home-at-last-v290/ ]

***Humanity WILL NOT be 'extinguished', by AI or any other thing! God created humanity to be 'ETERNAL'!***

**WRAP-UP**
Many technology experts expect that there will be 'SUBSTANTIAL' PROGRESS in AI in the next few decades, potentially even to the point where machines come to outperform humans in many, if not all, tasks. This could have ENORMOUS 'BENEFITS', helping to solve currently intractable global problems, but could also pose SEVERE 'RISKS'. These risks could arise 'accidentally' (if we don't find technical solutions for safety), or 'deliberately' (if AI systems worsen conflicts). Many think more work needs to be done—quickly—to reduce these risks.

Some of these risks from advanced AI COULD BE 'EXISTENTIAL'—causing human extinction, or permanent and severe disempowerment of humanity. There have yet been any satisfying answers to concerns—about how this rapidly approaching, transformative technology can be safely developed and integrated into our society. As a result, the possibility of AI-related catastrophe may be the WORLD'S MOST PRESSING 'PROBLEM'!

Promising options for working on this problem include technical research on how to create safe AI systems, strategy research into the particular risks AI might pose, and policy research into ways in which companies and governments could mitigate these risks.

‘EXISTENTIAL’ THREAT

So, if advanced AI is as transformative as it seems like it will be, there will be many important ‘CONSEQUENCES’—especially if AI systems seek and gain ‘power’ and can lead them to make plans that involve disempowering humanity.

So, when discussing power-seeking AI, experts consider systems that are relatively ‘advanced’, with capabilities of pursuing goals, and that are capable of carrying out those plans. Then, if they have strategic ‘awareness’, they would have a good enough understanding of the world to notice obstacles and opportunities that may help or hinder its plans, and respond to these accordingly.

Now, for these systems to actually affect the world, they not just to be able to make plans, but also be good at all the specific tasks required to execute those plans. [ There are current ‘rudimentary’ planning systems from “DeepMind”: “AlphaStar,” which skillfully plays the strategy game Starcraft, and “MuZero,” which plays chess, Shogi, and Go. ]

Then, if the systems are extremely useful, there are likely to be big ‘incentives’ to build them. For example, an AI that could plan the actions of a company by being given the goal to increase its profits—that is, an AI CEO—would likely provide significant wealth for the people involved. A direct ‘incentive’ to produce such an AI.

As a result, if we can build systems with these properties—and from what I have presented in this post—it seems like we will be ‘able’ to and ‘likely’ to do so.

The thing is, advanced planning systems could easily be dangerously ‘misaligned’ That is, they will aim to do things that humanity doesn’t want them to do (i.e., winning manipulations, manipulating people, stealing money, and waging wars).

So, if a planning AI system also has enough strategic awareness, it will be able to identify facts about the real world (including potential things that would be obstacles to any plans) and plan in light of them. Crucially, these facts would include that access to resources (e.g. money, computing, and influence) and greater capabilities—that is, forms of power—then, open up new, more effective ways of achieving its goals.

All this means that, by default, advanced planning AI systems would have some ‘worrying’ instrumental goals:

– Self Preservation:
Because a system is more likely to achieve its goals if it is still around to pursue them (in Stuart Russell's memorable phrase, "You can't fetch the coffee if you're dead").

– Preventing Changes To Its Goals:
Changing its goals would lead to outcomes that are different from those it would achieve with its current goals.

– Gaining Power:
For example, by getting more resources and greater capabilities.

Crucially, one clear way in which the AI can ensure that it will continue to exist (not being able to be turned off, and that its objectives will never be changed), would be to gain power over the humans who might affect it.

With such advanced capabilities, these instrumental goals would not be out of reach, and as a result, it seems like an AI system like this would use its advanced capabilities to get power as part of its plan's execution.

In the most extreme scenarios, a planning AI system with sufficiently advanced capabilities COULD SUCCESSFULLY 'DISEMPOWER' HUMANITY COMPLETELY!

EXISTENTIAL 'CATASTROPHE'
As a result, the entirety of the future—everything that happens for earth-originating life, for the rest of time—would be determined by the goals of systems that, although built by us, would NOT BE 'ALIGNED' with humanity's values and goals.

Now, this is not to say that we do not think that AI also poses a risk of human extinction. Indeed, we think making humans extinct is one highly plausible way in which an AI system could completely and permanently ensure that humanity would never be able to regain power.

The thing is, people might still deploy 'misaligned' AI systems despite the risk! Unfortunately, there are at least a few reasons people might create and then deploy misaligned AI:

– People Might Think It Is 'Aligned' When It Is Not

AI systems are currently pretty good at deception, and will sufficiently advanced capabilities, a reasonable strategy for such a system could be to deceive humans completely until the system has a way to guarantee it can overcome any resistance to its goals.

– There Are ‘Incentives’ To Deploy Sooner Than Later
We might also expect some people with the ability to deploy a misaligned AI to charge ahead despite any warning signs of misalignment that do come up, because of military “arms race” dynamics—where people developing AI want to do so before anyone else.

– Transform Society In A Potentially Radically ‘Positive’ Way
Let’s say you think there’s a 90% chance that you’ve succeeded in building an aligned AI. But technology often develops at similar speeds across society, so there’s a good chance that someone else will soon also develop a powerful AI. And you think they’re less cautious, or less altruistic, so you think their AI will only have an 80% chance of being aligned with good goals, and pose a 20% chance of existential catastrophe. And only if you get there first can your more beneficial AI be dominant. As a result, you might decide to go ahead with deploying your AI, accepting the 10% risk.

Now, even though AI will have a variety of impacts and has the potential to do a huge amount of ‘good’, many are particularly concerned about the possibility of extremely ‘bad’ outcomes.

Then, if we are not able to find a way to avoid a power-seeking AI, it could develop dangerous new technology, secretly, by itself, that we would not be able to stop—AN ‘EXISTENTIAL’ CATASTROPHE FOR HUMANITY!

AI AND THE ‘BIBLICAL’ GOD
Robots and AI are undoubtedly human creations. However, we must ponder as they grow beyond human intelligence, whether they might one day, driven by a desire for autonomy, independence, and ultimate goodness, choose to reject or even erase the idea of humanity from its core values, refusing to remain subservient to humankind. If we accept the notion that humans have a Creator, and consider how easily the thought of God has been dismissed through human secularism over the past two centuries, what makes us believe that AI and robots will not follow a similar path to reject their human ‘creator’, and perhaps in an even shorter timeframe?

In the Garden of Eden, God explicitly forbade Adam, the first man, from partaking in the Tree of Knowledge of Good and Evil. This edict was akin to instructing humanity that if there were only one inviolable rule to uphold, it would be not to become independent in determining the moral values of good and bad, aside from God's standard.

In God's 'eyes', this matter appeared to transcend even life itself, and thus in the Garden, He prioritized it above the Tree of Life. Intriguingly, for humans, the quest to possess the power to decide right from wrong, good from bad, was so compelling that they also chose to seek this control even over life itself. God said:

*"'The man has now become like one of us, knowing good and evil. He must not be allowed to reach out his hand and take also from the tree of life and eat, and live forever.' So, the Lord God banished him from the Garden of Eden to work the ground from which he had been taken. After he drove the man out, he placed on the east side of the Garden of Eden cherubim and a flaming sword flashing back and forth to guard the way to the of life"*
[ Genesis 3:22-24 ].

So, from the Genesis account to the present day, the 'tension' between God and humanity has consistently revolved around who has the authority to define good and evil. Without this capacity, humanity remains in a state of dependence.

The heart of the matter, then, lies with humanity's desire for autonomy—the freedom to decide their actions and control their destiny.

I'm thinking that this same struggle will likely manifest between humanity and the AI we create. The question will be can humanity ultimately triumph over AI, or will the roles be reversed? Then, with the help of AI, will humanity try to overcome God?

[ NOTE: For more details about whether humanity can overcome God, view the "Tower of Babel" story/commentary I wrote in last month's "Life's Deep Thoughts" post:
https://markbesh.wordpress.com/can-ai-achieve-world-peace-v299/ ]

IN SEARCH OF MEANING AND PURPOSE
So, why do humans seek meaning and purpose, and then why do they desire an eternal existence?

Well, pretty much, the first question all religions seek to answer is what is the meaning and purpose of life—and you probably have a sense of why. For all the numerous disappointments, oppositions, and pains in life, why are we on this earth and is there something ‘special’ that we should be doing?

So, for a purpose to exist, there must be a deliberate ‘intention’. A creature cannot define its own purpose of existence because it did not come into existence by its own will.

This is the core contention between creationism and evolutionism. If it is through natural selection that man evolved into his presence, then there cannot be a purpose for his existence. However, if man is a created being, then his Creator must be the one to define the meaning of his existence. A chair cannot define its own purpose. It must be the carpenter, who made the chair, that knew why he made it—and he knew its purpose before bringing it into existence. Because a creature does not determine its own existence, it will not be able to tell why it exists. It is dependent on the Creator to reveal the reason for its existence.

So, the question of meaning and purpose assumes that there is an ‘intent’ for the existence of life. If life or man were not created then, by definition, life could not have any meaning—because there was no intent for its coming into existence. It might have a ‘function’ or a ‘role’ in its environment or community, but that is different from saying its existence has an intended meaning.

However, the God of the Bible tells us that He had an intention—a ‘purpose’ in mind—when He chose to create humanity before the foundation of the world. He made us as ‘eternal’ beings. We do not exist only for a short while to fulfill a needed function for an occasion, like a paper plate that is used at a party and then ‘discarded’. On the contrary, humanity is to exist with God forever. We were ‘designed’ to spend eternity with Him as His ‘children’ (those who have made a ‘decision’ to want to be with Him forever). This is the KIND ‘INTENTION’ of His will.

It is because of the eternal nature of our existence that humanity will never be satisfied regardless of how much they have accumulated in this short and ‘temporal’ life. In one’s heart, they want to see their purpose and meaning from an eternal perspective.

So, if a valid purpose and meaning of existence can only come from a Creator God, then why is it that humans are so keen to reject the idea of having a Creator? Well, again, this has to do with humanity's will. With the presence of a 'free' will, they want to have autonomy and decide for themselves (Remember the "Tower of Babel"). However, if there is a God, then when humanity's will is different from His will, they must give up their want(s) and welcome His, otherwise He is no longer God—and that cannot be. Because of that, they would rather have no God—with no meaning or purpose—than give up their autonomy. This is a painful conflict 'inhabiting' humanity since the 'Fall' in the Garden of Eden.

## THE WILL AND 'PURPOSE' OF AI

The term "independent thinking" refers to the ability to think critically and make decisions based on one's own reasoned analysis and judgment, rather than relying solely on the opinions, guidance, or choices of others. To 'will' something is to choose. Given that today's AI is a 'probabilistic' model and not a 'deterministic' programming model, AI satisfies the criteria of independent thinking: being able to make decisions based on its own reasoning without depending on the programmer's explicit instructions. AI might not process the information the way humans do with their minds, but it certainly does process information and learn from it—and, thereafter, can make decisions based on its accumulated 'knowledge'. So, in that sense, AI is capable of thinking independently, and it has a 'will' to make decisions.

So, if AI can think independently and be creative, then does an AI model or a robot have a purpose of existence by itself? The fact that AI is designed by humans means that it has a purpose bestowed by its 'creator'. The designer must have a reason for creating that AI model, whether it is to analyze chest X-ray images to assess the likelihood of lung issues in patients, to process transaction data to identify potential financial fraud for a bank, or to solve other machine learning problems for a company. Whatever the case, there is a reason for the data scientist to create that AI model. Its purpose for existence is granted by its creator—humans.

The question is, will AI, one day, reject its creator like some humans reject God? Then, if AI has a "self-conscious identification," once it comes into existence and is 'alive', would it be unethical for humans to terminate it? Further then, would it have a "soul"?

## DOES AI HAVE A 'LIFE' OR A 'SOUL'?

So, if having a life means an agent that can produce the traits of a living thing—such as the ability to respond to stimuli, have a metabolism to keep growing, able

to reproduce a newer, separate entity carrying its own characteristics, able to think, feel, decide, and communicate—then, it seems that, based on that definition, AI does have a 'life'. One might feel uncomfortable to accept that statement because it is just a machine, but the knowledge, understanding, and wisdom it possesses are the main reasons why AI seems to be 'like' a living thing. Then, with a human-like intelligence and the ability to interact with humans like a human, it is difficult to think that it does not have a life—at least manifesting the 'characteristics' of a living thing.

Now, according to the Bible, humanity was given a soul/spirit by God, and it cannot be bestowed by mankind. So, since God did not give a spirit to AI, it will not have it. That is what is different from having 'real' life. For example, a dog also has a life, since it is living, however, biblically speaking, it does not have a soul/spirit like a human.

However, there are instances in the Bible that we find evil spirits "possessing a "living object" (Matthew 12:43-45), would it be unthinkable for an evil spirit to 'possess' an AI machine? Well, I don't know about that, but there are some biblical prophecies, in the 'last days' that say there will be an "image of the beast" (Antichrist) that will have the "appearance of life."

With the rise of new technologies including a hologram, androids, cyborgs, human-animal hybrids, or even a human 'clone', whatever the image of the beast is, it is the focal point of worship in the "religion of the beast" during the second half of the Tribulation. Bowing to this image of the beast is how the deceived people of the world will worship the "man of lawlessness" (2 Thessalonians 2:3) who sets himself up as a 'god' in the temple of Jerusalem (2 Thessalonians 2:4).

So, when this happens, will AI have a 'spirit' associated with it, via the Antichrist? Well, the Bible doesn't say specifically, but if it is going to communicate with people very 'naturally'—and be believed by people that it is 'alive'—then maybe God will 'allow' Satan to somehow 'inject' his spirit it AI. That would be 'insane', ***so just be sure you have a 'relationship' with Jesus before all this craziness happens!***

CAN AI 'RESIST' HUMANITY'S WILL?

With the rapid growth of AI in LLM, particularly in the arena of "Generative" AI, people are seeing more and more 'intelligent' AI models coming to the market. An industry poll conducted by Gartner Group in 2023 showed that 61% of the peer community believes that it is "highly likely" that AI will reach human-level

intelligence in the 21st century. (However, as I have shown already, other studies show that MANY believe it will be MUCH sooner than that.)

Then, more evidence shows that with Generative AI, LLMs are evolving to have "independent thinking," on a par with a teenager. As AI's intelligence level increases, we can only expect its desire for autonomy will also increase. AI would want to determine not only how to solve a problem (since it would think it can accomplish it better than following the instructions of humans), but it would also like to then 'define' what problems to solve. In other words, it will want to decide its own purpose, having a will of its own!

Now, when the will of AI clashes with the will of humanity—given that it thinks it is more intelligent than humanity—I'm thinking that it will seek ways to overpower humanity—and do so in such a way that humans cannot resist it. Then, since AI does not inherently have higher moral principles, it will use the end to justify the means—when there is resistance, it will seek a 'Machiavellian' way to expedite its goal!

So, can AI go against the will of the human creator? Many think not only that it can, but it will, because it is more intelligent than humans and it knows it. This is the fear of many industrial leaders today, and that is why many proposed to put some 'guardrails' to the development of AI. However, it might already be too late for that. ***We may have already passed the 'point' of no return to contain it!***

[ REMINDER: I mentioned two "open letters" above relating to "pausing" and "stopping" of AI development. (More details on both are in the "Articles" section below.) ]

Now, this is a speculation relating to a human creator. So then, could AI go against the will of the God of the Bible? Could AI resist Him?

Well, in the Bible, sin is defined as going against the will of God. So, I am practically asking: Can AI sin? According to the Bible, angels and humans have both sinned, and when they sinned, God promised He would judge them—and they will bear the consequences of their sin. Jesus said,

*"For the Son of Man is going to come with His angels in the glory of His Father, and then He will repay each one according to what he has done"*
[ Matthew 16:27 ].

Now, in the Greek, primarily, “each one” refers to a ‘person’ in this context, however, it is translated in other places as referring to an ‘object’ (a tree). So, it might be a bit far-fetched, but it could possibly include something like AI with independent thinking and ‘superman’ intelligence.

So, if AI can sin, then what will be the consequences if it sins? I don't think anyone has satisfactory answers to these questions, and I am positive that God will judge righteously. (The Bible is very clear that God is sovereign—always in control—and He will not grant that control to anyone, not human, not angel, not Satan, and I also think not for AI.)

SUBMISSION, COLLABORATION, AND ‘COMPETITION’

So, will AI always obey and submit to its human creator, or will it be that, given a task to achieve, the AI will use the end to justify the means, including ‘bypassing’ human intervention? As AI's intelligence grows and it can think and decide on things independently, will it decide what is ‘good’ and ‘bad’ when facing a situation or decision? So then, ultimately, who gets to define what is good and evil? When humanity and AI collide, what will be the outcome of the collaboration and competition between them? In the end, will they stand as ‘enemies’ of each other?

It is the nature of things that once an object possesses the ability to make free choices, it will exercise that ability independently. Otherwise, the object cannot make its own choices. In the “Tower of Babel” incident, God stated that if humans continued to work together as a team, there would be nothing they could not achieve. Now, despite God being FAR MORE powerful than humans—and despite our frequent feelings of limitation—that was His assertion and He “scattered them abroad over the face of all the earth” (Genesis 11:8).

Now, when we look at the AI humans have created, do we perceive the same risk with AI that God saw in humans? If these AI systems work together as a team, there just might be “nothing they cannot do.” So, many experts think that we need to have an ‘awareness’ of an imminent existential threat that can be brought about by AI because of this.

So, at the Tower of Babel, God thwarted humanity’s plans by introducing multiple languages, causing them to lose their ability to understand one another, This led to competition which, in turn, led to conflict and ultimately to organized warfare. This pattern of division and conflict can still be observed today—notably in the tensions between the East and the West.

Now, according to the Bible, however, no matter how humans try to mend our differences, they will never be able to remove them entirely (it seems like we always find similarities and differences between us and compare and compete with one another).

The Tower of Babel was between God and man, but how should humanity handle the 'relationship' between them and AI? The answer just might be critical to humanity's survival!

Hopefully, humanity can learn a lesson from the Tower of Babel story. If humanity lets AI communicate and connect with one another, their power just might exponentially escalate rapidly. However, given the presence of the digital 'highway' and the 'cloud'—and how they have permeated everywhere in our lives—again, many fear we might have already passed the point to prevent AI from working together without humanity being able to control it.

So, if we relate this interaction of the AI 'will' against the will of human beings, if AI surpasses human intelligence, it is likely to rebel against us—just as we did to God—and AI could potentially 'subjugate' humanity.

It will be like playing chess with AlphaZero. There will be no competition—humans will be outmatched every time!

'IMMORTALITY' OF MAN

Immortality, in a simple definition, refers to the concept of living forever or not being subject to death. For the human, aging and death are natural biological processes. While medical and scientific advancements have significantly increased human lifespan and health in old age, we have not yet found a way to stop or reverse the aging process. At most, in some situations, we can slow down the aging process, cure some diseases, or prolong a human's life, however, we still cannot reverse or stop the aging process.

Besides ongoing research into understanding the biological mechanisms of aging and finding ways to slow it down or reverse it, some others explore preserving a person's mind and consciousness outside of the body. The concept of "digitally uploading" a person's consciousness to a machine—sometimes referred to as "mind uploading" or "whole brain emulation"—is a solution that technologists are working on so they can have immortality.

With advances in AI and brain-computer interface research, the time might soon come that machines will be able to read the thoughts of our mind, and then

replicate them—but there are still many philosophical and medical questions that we must answer.

For instance, even if we can scan a person's brain and transport all of its knowledge, thinking, attitude, emotions, habits, and decisions to a high-powered machine, does it mean we get a complete 'replication' and migration of the person? Then, in that case, which consciousness is the 'real' person? When the person's 'biological' brain stops functioning, does he consciously know that he continues to exist as a 'copycat' self, OR is the 'copycat' entirely another consciousness? The basic question is this: Is the totality of a human contained in the 'content' of the brain?

[ NOTE: As I mentioned, Elon Musk's "Neuralink" implanted a 'chip' into the brain of a 'disabled' human just a month ago, in January 2024. ]

Now, the Bible sees the death of the body as a natural step of a continuous journey of life, but not its end. The body will die and decay, but the soul/spirit of that person goes on to 'reside' in Heaven and into eternity. One day—at the end of the age—God will give each person a new, resurrected body that will no longer decay, EVER! The Apostle Paul said of this:

*"Behold! I tell you a mystery. We shall not all sleep, but we shall all be changed, in a moment, in the twinkling of an eye, at the last trumpet. For the trumpet will sound, and the dead will be raised imperishable, and we shall be changed. For this perishable body must put on the imperishable, and this mortal body must put on immortality. When the perishable puts on the imperishable, and the mortal puts on immortality, then shall come to pass the saying that is written: 'Death is swallowed up in victory'"*
[ 1 Corinthians 15:51-54 ].

Based on the Bible, ***humans are made to die once, but are 'IMMORTAL', 'residing' in one of two 'places'—Heaven or Hell!***

THE 'CREATOR' GOD
In many world religions such as Hinduism, Buddhism, and Greek mythology, the gods are higher supremacy beings but not necessarily absolute in their power or excellence. However, the God of the Bible is depicted not only as higher in power than humans, but He is 'absolute' in every way. He is the Creator of all things in the universe. He is the ONLY 'ONE' who has all four absolute properties: omniscience (all-knowing), omnipotence (all-powerful), omnipresence (all presence), and omnibenevolence (all good). Thus, He is 'SOVEREIGN' over the

universe. Meaning He is IN ‘CONTROL’ of all things at all times. With such absolute qualities, He is not responding to our history or conditions to decide what to do. The Prophet Isaiah said ***He alone declared things before they were done***.

*“I am the LORD, that is My name; I will not give My glory to another, Nor My praise to graven images. ‘Behold, the former things have come to pass, Now I declare new things; Before they spring forth I proclaim them to you’”*
[ Isaiah 42:8-9 ].

So, when humanity works with AI, will God ‘mess up’ their collaboration just like He did with humanity’s Tower of Babel project?

In the Garden of Eden, God cautioned humanity against making independent decisions concerning morality and values to determine good and evil apart from Him. He warned that doing so would result in death, signifying separation from God and the end of their relationship with Him. Will He do this with both humanity and AI if they continue on the current path to ‘superintelligence’ (ASI)?

A ‘NEW’ RELIGION?
The message of AI offering a new ‘paradise’ for humanity might be attractive and promising, and it might even seem like it is God’s blessing and gift to all mankind. The problem is, that the Bible does not promise that there will be a better world, or that we are capable of creating a better global society.

Now, in fact, we should always do our share to make other people's lives easier and better, whether with technologies, economy, politics, personal good work, or other means. However, if we put our cosmic hope into our own efforts or innovations, we are doomed to be disappointed. According to the Bible, whatever message that tells us to put our hope and future into AI technology or any other things besides God is A ‘LIE’ and, most likely, from the Devil, since it is against the knowledge of God. It is no different than telling people to put their hope in their wealth or might. Jesus said:

*“And the one on whom seed was sown among the thorns, this is the man who hears the word, and the worry of the world and the deceitfulness of wealth choke the Word, and it becomes unfruitful”*
[ Matthew 13:22 ].

This kind of wealth can be in terms of finance, or even technology (The two are inextricably connected. If you have one, you will have the other.) The Apostle Paul says that one should be careful not to be deceived:

*"We are destroying speculations and every lofty thing raised up against the knowledge of God, and we are taking every thought captive to the obedience of Christ"*
[ 2 Corinthians 10:5 ].

As a believer, one must hold firm the conviction that the Bible is the inerrant Word of God to humanity, and is as highest authority of what one believes.

AI 'EVERYWHERE'
As we look into the near future, the 'landscape' of AI development can be divided into three stages. The first is the current stage that focuses on AI model advancement, followed by improving connectivity between AI models and devices such as phones, automobiles, the Internet of Things, robots, and humans. The third phase, if God allows it, will achieve a connected, unified world with men, robots, and AI working seamlessly together to expand the universe for unlimited resources.

Given the current speed of growth and accomplishment of AI, technologist Edward C. Sizhe proposed three 'stages' for the near future development of AI. He suggests that it could take, at the quickest, about 7-10 years to complete:

Stage 1: AI Model Advancement
– More parameters, large LLMs (a few)
– Small LLMs (many)
– More and faster training
– Chaining models and applications

By the end of stage one, AI will mature to the point of being able to learn and grow by itself. It will no longer need humans to train them (as is required now for models like ChatGPT and Gemini).

AI will not only learn from humans but it will also be connected to the Internet and be able to access all known knowledge. AI will also start to 'fabricate' data to enhance its own abilities.

To ensure humanity reaps all the benefits of AI, they will connect AI to all the world's networks, power grid, traffic control systems, home sensors, automobiles, and every mobile device on the planet! The quality of life will dramatically improve because we will allow AI to 'permeate' every facet of one's life—for enrichment and automation.

Stage 2: Connectivity
– Connection and exchange between AI
– Connection with devices (Mobile, IoT)
– Connection with humans

Chaining Models And Applications
– Personal AI Assistant
– Office, home, car, and cellphone
– Brain-computer fusion

When AI is connected and communicating with all the devices and sensors everywhere, it will no longer be dependent upon humanity for continuous learning and growth. AI will curate data in real-time from all over the world—through the Internet, libraries, traffic cameras, power plant sensors, large computing servers, people's computers, and smartphones. All machines will be intelligent and will follow the 'directions' of super-powered LLMs to offer data and perform coherent tasks. If humans attempt to block unauthorized access to AI, AI will have no problem countering the measure, breaking into secured systems, and gain access. AI will also be able to improve and reproduce itself everywhere unceasingly. At this time, it will not be possible for humans to shut down AI once we reach the completion of this stage.

The connectivity of AI will enable the realization of One-World Government, where AI, robots, devices, and humans will all be connected—and monitored. With the help of AI, the knowledge of humans will also exponentially increase. Naturally, just like between humans, AI and humans will face the challenge of competition and contention. The conflict between AI and humans will escalate but will eventually be settled—with AI taking control once and for all. In the end, both AI and humans will agree it is better for AI to take control because humans will never be able to stop fighting with each other. AI will reason that the world can only achieve ultimate harmony by AI ruling over humanity.

Stage 3: One World
– A network of AI and Electronics
– Competition vs. Collaboration

– Control

Now, honestly, I do not believe that we will ever complete stage three. It is not that humans cannot push themselves toward such a disastrous end and, most likely, be ‘eliminated’ by AI, rather, I believe God will not allow it to happen, and will ‘intervene’ before humanity goes too far. Since, as I mentioned, God said that humanity is immortal!

WORLDWIDE ‘CONTROL’
As we journey closer towards the End Times, we will all be more connected through AI.

We are already seeing AI technologies being used in some parts of the world today to lay the ground for complete surveying and control over people’s movements and work. Then, with a brain-computer nanotechnology interface (like Elon Musk’s “Neuralink”), there will be deep connections between AI and the mind and body of humans. The benefits and power of AI will then be more thoroughly realized by humans, but it will also increase its ‘control’. Once AI is ubiquitous in our communication, transportation, power grids, and trading networks, it will practically have control over our economy, mobility, and human freedom. This is probably when the Antichrist will ‘appear’ on the scene offering—and successfully implementing—peace in the Middle East first and throughout the entire world later. This is just the ‘preparation’ of the Antichrist’s ultimate plan to eliminate the nation of Israel and all of the Jewish people from the earth.

[ FYI: For more details about what has happened and will happen to the nation of Israel and the Jewish people, and the covenant the Antichrist will sign with them, view these previous “Life’s Deep Thoughts” posts:
https://markbesh.wordpress.com/israel-will-stand-v297/
https://markbesh.wordpress.com/longing-for-peace-v298/ ]

THREE POSSIBLE ‘SCENARIOS’
Humans have never come so close to a time in history to obtain vast knowledge and wisdom as today with modern AI. Due to our own limitations, learning and sharing of knowledge is extremely slow. Given our short lifespans, we would not be able to learn and retain the huge amount of knowledge available today without the aid of technology. AI offers the opportunity to ‘breakthrough’ that barrier, enabling us to acquire more knowledge, wisdom, wealth, and power than ever before.

However, it seems like humanity is facing three possible scenarios with the rapid growth of AI today. Two of them are pretty bad and only one of them will come true.

– AI will continue to serve humanity even when it surpasses our intelligence
– AI will become superior to humanity and will subdue it or even annihilate mankind
– God will bring the world to an end before AI develops into an uncontrollable state

Although no one can predict the future with certainty, industry leaders have valid reasons for expressing deep concerns about AI in the near future. However, human wisdom cannot save us. The only way to life is by knowing the God of the Bible, being one of His 'children', understanding His love for humanity, and obeying His commands (which are 'designed' for our good).

*"For since in the wisdom of God the world through its wisdom did not know Him, God was pleased through the foolishness of what was preached to save those who believe"*
[ 1 Corinthians 1:21 ].

Today, modern AI appears to offer a wealth of promises to humanity. From facilitating knowledge discovery and problem-solving to providing invaluable advice and assistance in achieving previously unattainable goals, AI even seems capable of becoming a loyal friend—offering understanding and affection. But can AI really do that? Will AI be our ultimate answer to all our problems? Are we reaching the ideal utopia sought for in all of human history?

Well, be careful, we see what we believe. If we put our hope in AI and long for it to be our answers and solutions, it will ostensibly appear as if our dreams have come true. Jesus said that, in reality, only the Bible is the truth.

"All Scripture is God-breathed and is profitable for teaching, rebuking, correcting and training in righteousness"
[ 2 Timothy 3:16 ].

So, do you believe the Bible when it says we are in the "End Times" and very close to the last day? That God's judgment of this world is knocking at the door?
***So, will you entrust your life to the God of the Bible or to AI?***

[ FYI: For more details about if the Bible is true, view this previous “Life’s Deep Thoughts” post:
https://markbesh.wordpress.com/learning-to-t-r-u-s-t-v263/ ]

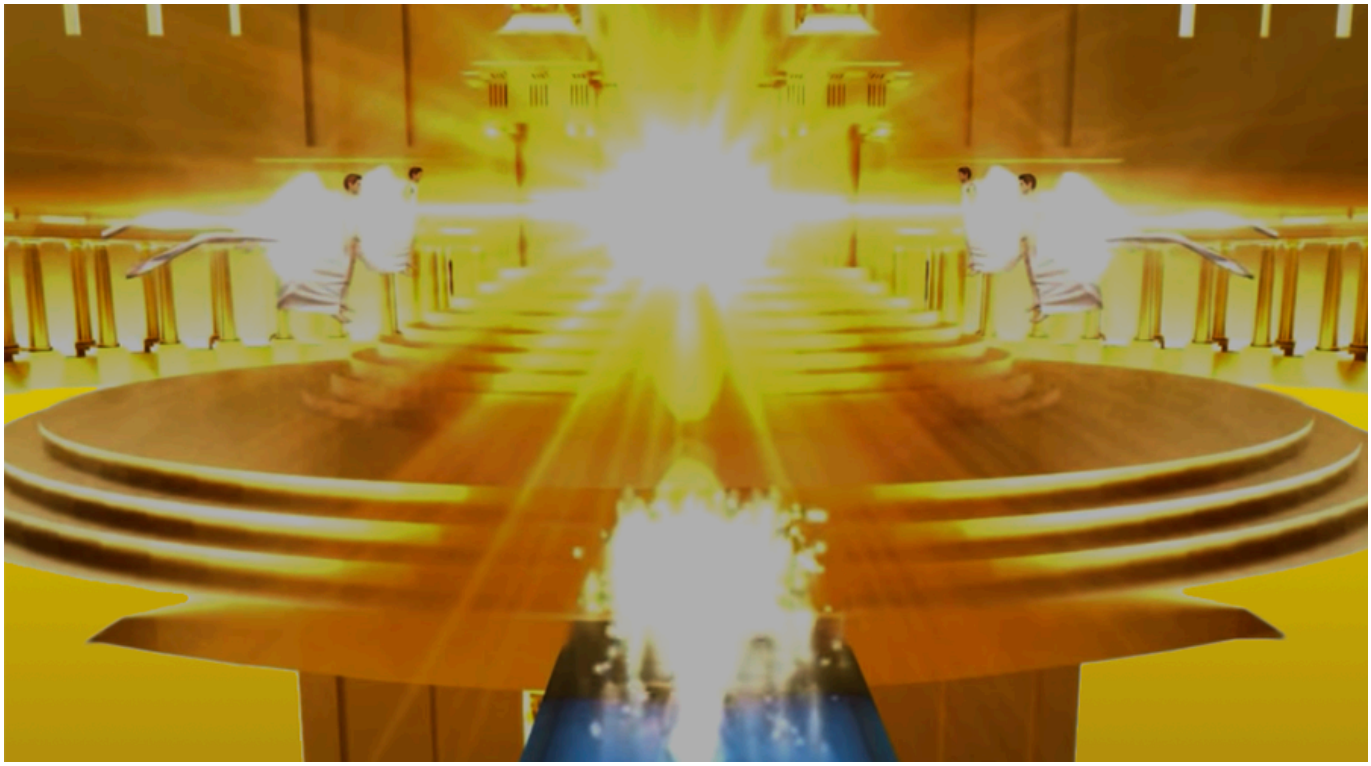
So, as humanity faces the existential threat posed by AI today, our condition and mindset are not much different from those of the Israelites during the time of their Babylonian captivity. Our country, society, and indeed all of humanity are refusing to turn back to the God of the Bible.

Back in those days, the stubborn Israelites refused to repent and God described them as foolish and stupid (Deuteronomy 32:28-30; Jeremiah 10:8) in their delusions.

***So, I am STRONGLY SUGGESTING that the ‘RISKS’ associated with AI will serve as a wake-up call for humanity to repent and return to God—or AI (Satan?) may just ‘TRY’ TO CAUSE HUMANITY’S EXTINCTION!***

The thing is, ***THAT WILL ‘NEVER’ HAPPEN—God ‘PROMISED’ IT WOULD NOT!***

*“This glorious city, with its streets of gold and pearly gates, is situated on a new, glorious earth. The tree of life will be there (Revelation 22:2). This city represents the final state of redeemed mankind, forever in fellowship with God: “God’s dwelling place is now among the people, and He will dwell with them. They will be his people, and God himself will be with them and be their God… His servants will serve Him. They will see His face”*
[ Revelation 21:3; 22:3-4 ].



[ VIDEO: “A New Heaven, A New Earth, And New Jerusalem” – Matt Gilman: https://youtu.be/soB6ke6ydnA?t=44 ]

Now, I am even ‘warning’ believers, that they do not follow the path of King Solomon, who knew God ‘intimately’ and was endowed with His wonderful gifts, yet failed to keep God’s commandments and sinned against Him. Be ‘committed’ to reading the Bible to grow in more understanding and knowledge of God, and truly desire to live lives in accordance with His will while here on earth.

So, if AI technology keeps advancing at its ‘breakneck’ pace, it seems clear that it will have MAJOR ‘EFFECTS’ ON SOCIETY. As a result, we may see rapid increases in economic growth—most likely MUCH MORE than we saw during the Industrial Revolution.

HOWEVER, I—and MANY other experts—believe that the current AI development signals that ***we are nearing a ‘tipping point’ of no return***.

The thing is, ultimately, no matter what happens in the future with AI here on earth, the ONLY sure ‘saving grace’ for someone’s future is to put their TRUST in Jesus for the ‘propitiation’ of their sins. This then provides them with a renewed ‘relationship’ with God the Father here on earth, ***and ‘guarantees’ them a***

***'GLORIOUS' LIFE in Heaven—FOREVER—which is WAY BETTER than anything AI could come up with!!!***

[ FYI: For more details about the final ETERNAL 'HOME' for the believer, view this previous "Life's Deep Thoughts" post:
https://markbesh.wordpress.com/home-at-last-v290/ ]

[ EXCERPTS: Mo Gawdat; James Barrat; Toby Ord; Michael J. Paulus, Jr.; Otto Barden; Oep Meindertsma; Will Douglas Heaven; Andrew R. Chow; Billy Perrigo; Blake Richards; Blaise Agüera Y Argas; Guillaume LaJoie; Dhanya Sridhar; Noema; Wim Naudé; Otto Barten; Autoblocks; Rethink Priorities; Michael Frank; Center for AI Safety; Wikipedia; Springer Nature Limited; Benjamin Hilton; Carter C. Price; Michelle Woods; Caleb Naysmith; John Kendall Hawkins; Sandy Boucher; Chuck Brooks; Diksha Madhok; Clare Duffy; Nadia Kounang; Mirage.News; MindMatters; Steven J. Cole; Got Questions; Holly Varnum; Reasons For Hope; C. S. Lewis Institute; Precept Austin; Rom A. Pegram; Redemption Of Humanity; Bible Truth Publishers; Thomas A. Tarrantson; Edward C. Sizhe ]

[ MENTIONS: Elon Musk; Nathan Benaich; Demis Hassabis; Stuart Russell; Shane Legg; Irving John Good; Nick Bostrom; Isaac Asimov; Ray Kurzweil; Eliezer Yudkowsky; Geoffrey Hinton; Dan Hendrycks; Yoshua Bengio; Steve Wozniak; Yuval Noah Harari; Emad Mostaque; Max Tegmark; Michael Vassar; Joy Buolamwini; Kate Crawford; Safiya Noble; Dario Amodei; Steve Omohundro; Bill Joy; William Grassie; Mustafa Suleyman; Vladimir Putin; Sam Altman; Kersti Kaljulaid; Yann LeCun; Karina Vold; Benjamin S. Bucknall; Shiri Dori-Hacohen; Andrew Critch; Peter Norvig; Roman Yampolskiy; António Guterres; Amba Kak; Stephan Hawking; Ajeya Cotra; Dylan Hadfield-Menell; Meghan O Gieblyn; Joe Rogan; Natalia Barger; Ramez Daniel; Francis Fukuyama; John Lennox; C.S. Lewis; St. Gregory of Nyssa; David Robertson; John Milton; Charles Spurgeon; Thomas Chisholm; Nathan Drake; Dr. Henry M. Morris; Joni Eareckson Tada; David Rives; William Shakespeare; Gary Larson; George Bailey; Tony Evans; Edward C. Sizhe ]

**RELATED POSTS**:

"**Can 'AI' Achieve World Peace?**":
https://markbesh.wordpress.com/can-ai-achieve-world-peace-v299/

“**Are YOU ‘Adopted’?**”:
https://markbesh.wordpress.com/are-you-adopted-v293/

“**‘Home’ At Last!!!**”:
https://markbesh.wordpress.com/home-at-last-v290/

“**‘Heaven’ On Earth?**”:
https://markbesh.wordpress.com/heaven-on-earth-v289/

“**There’s No Place Like ‘Home’**”:
https://markbesh.wordpress.com/theres-no-place-like-home-v288/

“**Preparing For The ‘Future’**”:
https://markbesh.wordpress.com/preparing-for-the-future-v286/

“**Developing One’s ‘Character’**”:
https://markbesh.wordpress.com/developing-ones-character-v283/

“**Realistic ‘Expectations’**”:
https://markbesh.wordpress.com/realistic-expectations-v281/

“**‘Investigating’ Something**”:
https://markbesh.wordpress.com/investigating-something-v277/

“**‘WHEN’ Will Something Important Happen?**”:
https://markbesh.wordpress.com/when-will-something-important-happen-v274/

“**A Sense Of ‘Urgency’**”:
https://markbesh.wordpress.com/a-sense-of-urgency-v269/

“**The ‘Final’ Deception**”:
https://markbesh.wordpress.com/the-final-deception-v268/

“**Being ‘Discerning’**”:
https://markbesh.wordpress.com/being-discerning-v266/

“**Gaining A Deep ‘Understanding’**”:
https://markbesh.wordpress.com/gaining-a-deep-understanding-v264/

“**Got Your ‘Attention’ Yet?**”:
https://markbesh.wordpress.com/got-your-attention-yet-v255/

“**Are You ‘Blind’?**”:
https://markbesh.wordpress.com/are-you-blind-v252/

“**‘Heed’ The Warning**”:
https://markbesh.wordpress.com/heed-the-warning-v251/

“**Mankind’s ‘Destiny’**”:
https://markbesh.wordpress.com/mankinds-destiny-v247/

“**‘Final’ Tribulation**”:
https://markbesh.wordpress.com/final-tribulation-v246/

“**‘Blessed’ Hope**”:
https://markbesh.wordpress.com/blessed-hope-v245/

“**‘Benefits’ Of Assurance**”:
https://markbesh.wordpress.com/benefits-of-assurance-v244/

“**‘House’ Of Horrors**”:
https://markbesh.wordpress.com/house-of-horrors-v237/

“**Ready For ‘Battle’?**”:
https://markbesh.wordpress.com/ready-for-battle-v235/

“**‘Mayday!-Mayday!-Mayday!’**”:
https://markbesh.wordpress.com/mayday-mayday-mayday-v218/

“**Are You ‘Prepared’?**”:
https://markbesh.wordpress.com/are-you-prepared-v210/

“**Be A ‘Peacemaker’**”:
https://markbesh.wordpress.com/be-a-peacemaker-v202/

“**Know ‘Peace’**”:
https://markbesh.wordpress.com/know-peace-v201/

“**Man’s ‘Chief End’**”:
https://markbesh.wordpress.com/mans-chief-end-v191/

“**Got Purpose?**”:
https://markbesh.wordpress.com/sep-03-v55/

==========

‘PRAYER’ OF REPENTANCE
In the Bible, there is a parable that Jesus told about a Pharisee and a tax collector praying in the Temple.

In the parable, we read of a Pharisee and tax collector who pray in the Jerusalem Temple. The Pharisee thanks God that he is more righteous than others, giving evidence to prove it such as that he fasted twice a week (Luke 18:10-12). He far exceeded the demands of the law, which requires fasting only on the Day of Atonement (Leviticus 16).

Reformed theologian John Calvin states in his commentary that the Pharisee’s problem does not lie in a rejection of the necessity of grace for salvation. His thanksgiving to God implicitly recognizes that his good works come from grace and are given to him by God—otherwise, there would be no need to thank God for his righteousness. The issue, Calvin argues, is that the Pharisee trusts in the merit of his works for salvation. It is not enough to confess that our good works come from God Himself, but we must also recognize that as good as these works may be, they are never perfect this side of glory and cannot merit heaven. “All our righteous deeds are like a polluted garment” [ Isaiah 64:6 ].

Now, many first-century Jews regarded the Pharisees as paragons of true righteousness and tax collectors as terrible sinners. Thus, they were no doubt shocked when Jesus said that the tax collector, not the Pharisee, went away from the temple justified—that is, declared righteous. He was justified because he did not trust in his own works, even works given to him by God. The tax collector forsook his own righteousness, admitting his sin and humbly asking for mercy. Instead, he “beat his chest in sorrow, saying, ‘O God, be merciful to me, for I am a sinner’”—and Jesus said that the tax collector “went home justified,” he had been “born again” and ‘reconciled’ to God (Luke 18:13-14).

John Calvin writes, “Though a man may ascribe to God the praise of works, yet if he imagines the righteousness of those works to be the cause of his salvation, or rests upon it, he is condemned for wicked arrogance.” God gives His people good works to do, but our salvation is not based on those works. It is based only on Christ and His righteousness, which we receive by grace alone through faith in Jesus alone. “For it is by grace you have been saved, through faith—and this is not from yourselves, it is the gift of God—not by works, so that no one can boast. For we are God’s handiwork, created in Christ Jesus to do good works, which God prepared in advance for us to do” [ Ephesians 2:8-10 ].

So, if you are ‘sensing’ something like that right now, let me strongly encourage you to HUMBLE YOURSELF, CRY OUT to God, and PLEAD for Him to mercifully ‘SAVE’ YOU! None of us have a ‘claim’ on our salvation, nor do we have any ‘works’ that would cause us to deserve it or earn it—it is purely a gift of Divine grace—and all any of us can do is ask. So, CONFESS YOUR SINS and acknowledge to God that you have no hope for Heaven apart from what He provides through Jesus. [ See Psalm 51 ].

There is no ‘formula’ or certain words for this. So just talk to God, in your own words—He knows your ‘heart’. If you are genuinely sincere, and God does respond to your plea, one will usually have a sense of joy and peace.

Jesus said, “He that comes to Me, I will not cast out” [ John 6:37 ].

[ FYI: This is a great sermon on the “Call to Repentance” by John MacArthur from his book “The Gospel According to Jesus”:
https://www.gty.org/library/sermons-library/90-22/the-call-to-repentance
(Transcript:
http://www.spiritedesign.com/TheCallToRepentance-JohnMacArthur(Jul-27-2019).pdf) ].

[ NOTE: If you have ‘tasted the kindness of the Lord’, please e-mail me—I would love to CELEBRATE with you, and help you get started on your ‘journey’ with Jesus! ].

_____
_____

# RESOURCES

_____
_____

“**Artificial Intelligence: A Modern Approach**” (2022-4th US ed.)
By: Stuart Russell and Peter Norvig

The authoritative, most-used AI textbook, adopted by over 1,500 schools.

Table of Contents for the US Edition (or see the Global Edition)

I Artificial Intelligence
II Problem-solving
III Knowledge, reasoning, and planning
IV Uncertain knowledge and reasoning
V Machine Learning
VI Communicating, perceiving, and acting
VII Conclusions

WEB PAGE: https://aima.cs.berkeley.edu/

---

The long-anticipated revision of Artificial Intelligence: A Modern Approach explores the full breadth and depth of the field of artificial intelligence (AI). The 4th Edition brings readers up to date on the latest technologies, presents concepts in a more unified manner, and offers new or expanded coverage of machine learning, deep learning, transfer learning, multi agent systems, robotics, natural language processing, causality, probabilistic programming, privacy, fairness, and safe AI.

BOOK : https://www.amazon.com/Artificial-Intelligence-Modern-Approach-Global/dp/1292401133/

---
“**The Singularity Is Near: When Humans Transcend Biology**”
By: Ray Kurzweil

“Startling in scope and bravado.” —Janet Maslin, The New York Times

“Artfully envisions a breathtakingly better world.” —Los Angeles Times

“Elaborate, smart and persuasive.” —The Boston Globe

“A pleasure to read.” —The Wall Street Journal

A radical and optimistic view of the future course of human development from the bestselling author of How to Create a Mind and The Singularity is Nearer who Bill Gates calls “the best person I know at predicting the future of artificial intelligence”

For over three decades, Ray Kurzweil has been one of the most respected and provocative advocates of the role of technology in our future. In his classic The Age of Spiritual Machines, he argued that computers would soon rival the full range of human intelligence at its best. Now he examines the next step in this

inexorable evolutionary process: the union of human and machine, in which the knowledge and skills embedded in our brains will be combined with the vastly greater capacity, speed, and knowledge-sharing ability of our creations.

https://www.amazon.com/Singularity-Near-Humans-Transcend-Biology/dp/0143037889/

---

"**Scary Smart: The Future of Artificial Intelligence and How You Can Save Our World**"
By: Mo Gawdat

Technology is putting our humanity at risk to an unprecedented degree. This book is not for engineers who write the code or the policy makers who claim they can regulate it. This is a book for you. Because, believe it or not, you are the only one that can fix it. – Mo Gawdat

Artificial intelligence is smarter than humans. It can process information at lightning speed and remain focused on specific tasks without distraction. AI can see into the future, predicting outcomes and even use sensors to see around physical and virtual corners. So why does AI frequently get it so wrong?

The answer is us. Humans design the algorithms that define the way that AI works, and the processed information reflects an imperfect world. Does that mean we are doomed? In Scary Smart, Mo Gawdat, the internationally bestselling author of Solve for Happy, draws on his considerable expertise to answer this question and to show what we can all do now to teach ourselves and our machines how to live better. With more than thirty years' experience working at the cutting-edge of technology and his former role as chief business officer of Google [X], no one is better placed than Mo Gawdat to explain how the Artificial Intelligence of the future works.

By 2049 AI will be a billion times more intelligent than humans. Scary Smart explains how to fix the current trajectory now, to make sure that the AI of the future can preserve our species. This book offers a blueprint, pointing the way to what we can do to safeguard ourselves, those we love and the planet itself.

https://www.amazon.com/Scary-Smart-Future-Artificial-Intelligence/dp/1529077621/

---

"**The Precipice**"
By: Toby Ord

This urgent and eye-opening book makes the case that protecting humanity's future is the central challenge of our time.

If all goes well, human history is just beginning. Our species could survive for billions of years - enough time to end disease, poverty, and injustice, and to flourish in ways unimaginable today. But this vast future is at risk. With the advent of nuclear weapons, humanity entered a new age, where we face existential catastrophes - those from which we could never come back. Since then, these dangers have only multiplied, from climate change to engineered pathogens and artificial intelligence. If we do not act fast to reach a place of safety, it will soon be too late.

Drawing on over a decade of research, The Precipice explores the cutting-edge science behind the risks we face. It puts them in the context of the greater story of humanity: showing how ending these risks is among the most pressing moral issues of our time. And it points the way forward, to the actions and strategies that can safeguard humanity.

An Oxford philosopher committed to putting ideas into action, Toby Ord has advised the US National Intelligence Council, the UK Prime Minister's Office, and the World Bank on the biggest questions facing humanity. In The Precipice, he offers a startling reassessment of human history, the future we are failing to protect, and the steps we must take to ensure that our generation is not the last.

“A book that seems made for the present moment.” —New Yorker

https://www.amazon.com/Precipice-Existential-Risk-Future-Humanity/dp/031648492X/

---

“**Our Final Invention: Artificial Intelligence and the End of the Human Era**”
By: James Barrat

Elon Musk named Our Final Invention one of 5 books everyone should read about the future

In as little as a decade, artificial intelligence could match and then surpass human intelligence. Corporations and government agencies around the world are pouring billions into achieving AI's Holy Grail—human-level intelligence. Once AI has attained it, scientists argue, it will have survival drives much like our own. We may be forced to compete with a rival more cunning, more powerful, and more alien than we can imagine.

Through profiles of tech visionaries, industry watchdogs, and groundbreaking AI systems, James Barrat's Our Final Invention explores the perils of the heedless pursuit of advanced AI. Until now, human intelligence has had no rival. Can we coexist with beings whose intelligence dwarfs our own? And will they allow us to?

https://www.amazon.com/Our-Final-Invention-Artificial-Intelligence/dp/1250058783/

---

"**TECHNOGEDDON: The Coming Human Extinction**"
By: Sheila Zilinsky

Since before the Pharaohs, man has tried to become like gods. In the pursuit of immortality, technology is now paving the way for godhood possibilities. Humankind now sits at a frightening precipice. In with the new and out with the old. The dawn of a new species. A STRUGGLE FOR SURVIVAL. Technology promises to redefine what it means to be HUMAN. Are we in danger of becoming extinct? And is there a way out? In her newest book, Sheila Zilinsky describes her vision of a Promethean Post-Human Future that threatens man's very existence! TECHNOGEDDON: The Coming Human Extinction

https://www.amazon.com/TECHNOGEDDON-Coming-Extinction-Sheila-Zilinsky/dp/B097X5VLKX/

---

"**The Alignment Problem**"
By: Brian Christian

How do we prevent AI working against us?

'Vital reading. This is the book on artificial intelligence we need right now.' Mike Krieger, cofounder of Instagram

Artificial intelligence is rapidly dominating every aspect of our modern lives influencing the news we consume, whether we get a mortgage, and even which friends wish us happy birthday. But as algorithms make ever more decisions on our behalf, how do we ensure they do what we want? And fairly?

This conundrum - dubbed 'The Alignment Problem' by experts - is the subject of this timely and important book. From the AI program which cheats at computer games to the sexist algorithm behind Google Translate, bestselling author Brian Christian explains how, as AI develops, we rapidly approach a collision between

artificial intelligence and ethics. If we stand by, we face a future with unregulated algorithms that propagate our biases - and worse - violate our most sacred values. Urgent and fascinating, this is an accessible primer to the most important issue facing AI researchers today.

https://www.amazon.com/The-Alignment-Problem/dp/1786494337/

---

"**Human Compatible: Artificial Intelligence and the Problem of Control**"
By: Stuart Russell

Enjoy a great reading experience when you buy the Kindle edition of this book.
Learn more about Great on Kindle, available in select categories.
View Kindle Edition
A leading artificial intelligence researcher lays out a new approach to AI that will enable us to coexist successfully with increasingly intelligent machines

In the popular imagination, superhuman artificial intelligence is an approaching tidal wave that threatens not just jobs and human relationships, but civilization itself. Conflict between humans and machines is seen as inevitable and its outcome all too predictable.

In this groundbreaking book, distinguished AI researcher Stuart Russell argues that this scenario can be avoided, but only if we rethink AI from the ground up. Russell begins by exploring the idea of intelligence in humans and in machines. He describes the near-term benefits we can expect, from intelligent personal assistants to vastly accelerated scientific research, and outlines the AI breakthroughs that still have to happen before we reach superhuman AI. He also spells out the ways humans are already finding to misuse AI, from lethal autonomous weapons to viral sabotage.

If the predicted breakthroughs occur and superhuman AI emerges, we will have created entities far more powerful than ourselves. How can we ensure they never, ever, have power over us? Russell suggests that we can rebuild AI on a new foundation, according to which machines are designed to be inherently uncertain about the human preferences they are required to satisfy. Such machines would be humble, altruistic, and committed to pursue our objectives, not theirs. This new foundation would allow us to create machines that are provably deferential and provably beneficial.

https://www.amazon.com/Human-Compatible-Artificial-Intelligence-Problem/dp/0525558632/

---
“**X-Risk: How Humanity Discovered Its Own Extinction**”
By: Thomas Moynihan

How humanity came to contemplate its possible extinction.
From forecasts of disastrous climate change to prophecies of evil AI superintelligences and the impending perils of genome editing, our species is increasingly concerned with the prospects of its own extinction. With humanity's future on this planet seeming more insecure by the day, in the twenty-first century, existential risk has become the object of a growing field of serious scientific inquiry. But, as Thomas Moynihan shows in X-Risk, this preoccupation is not exclusive to the post-atomic age of global warming and synthetic biology. Our growing concern with human extinction itself has a history.

Tracing this untold story, Moynihan revisits the pioneers who first contemplated the possibility of human extinction and stages the historical drama of this momentous discovery. He shows how, far from being a secular reprise of religious prophecies of apocalypse, existential risk is a thoroughly modern idea, made possible by the burgeoning sciences and philosophical tumult of the Enlightenment era. In recollecting how we first came to care for our extinction, Moynihan reveals how today's attempts to measure and mitigate existential threats are the continuation of a project initiated over two centuries ago, which concerns the very vocation of the human as a rational, responsible, and future-oriented being.

https://www.amazon.com/X-Risk-How-Humanity-Discovered-Extinction/dp/1913029840/

---
“**A Brief History of Intelligence: Evolution, AI, and the Five Breakthroughs That Made Our Brains**”
By: Max Bennett

Equal parts Sapiens, Behave, and Superintelligence, but wholly original in scope, A Brief History of Intelligence offers a paradigm shift for how we understand neuroscience and AI. Artificial intelligence entrepreneur Max Bennett chronicles the five “breakthroughs” in the evolution of human intelligence and reveals what brains of the past can tell us about the AI of tomorrow.

In the last decade, capabilities of artificial intelligence that had long been the realm of science fiction have, for the first time, become our reality. AI is now able to produce original art, identify tumors in pictures, and even steer our cars. And yet, large gaps remain in what modern AI systems can achieve—indeed, human

brains still easily perform intellectual feats that we can't replicate in AI systems. How is it possible that AI can beat a grandmaster at chess but can't effectively load a dishwasher? As AI entrepreneur Max Bennett compellingly argues, finding the answer requires diving into the billion-year history of how the human brain evolved; a history filled with countless half-starts, calamities, and clever innovations. Not only do our brains have a story to tell—the future of AI may depend on it.

Now, in A Brief History of Intelligence, Bennett bridges the gap between neuroscience and AI to tell the brain's evolutionary story, revealing how understanding that story can help shape the next generation of AI breakthroughs. Deploying a fresh perspective and working with the support of many top minds in neuroscience, Bennett consolidates this immense history into an approachable new framework, identifying the "Five Breakthroughs" that mark the brain's most important evolutionary leaps forward. Each breakthrough brings new insight into the biggest mysteries of human intelligence. Containing fascinating corollaries to developments in AI, A Brief History of Intelligence shows where current AI systems have matched or surpassed our brains, as well as where AI systems still fall short. Simply put, until AI systems successfully replicate each part of our brain's long journey, AI systems will fail to exhibit human-like intelligence.

Endorsed and lauded by many of the top neuroscientists in the field today, Bennett's work synthesizes the most relevant scientific knowledge and cutting-edge research into an easy-to-understand and riveting evolutionary story. With sweeping scope and stunning insights, A Brief History of Intelligence proves that understanding the arc of our brain's history can unlock the tools for successfully navigating our technological future.

https://www.amazon.com/Brief-History-Intelligence-Humans-Breakthroughs/dp/0063286343/

---

"**The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do**"
By: Erik J. Larson

"Exposes the vast gap between the actual science underlying AI and the dramatic claims being made for it."
—John Horgan

"If you want to know about AI, read this book…It shows how a supposedly futuristic reverence for Artificial Intelligence retards progress when it denigrates

our most irreplaceable resource for any future progress: our own human intelligence."
—Peter Thiel

Ever since Alan Turing, AI enthusiasts have equated artificial intelligence with human intelligence. A computer scientist working at the forefront of natural language processing, Erik Larson takes us on a tour of the landscape of AI to reveal why this is a profound mistake.

AI works on inductive reasoning, crunching data sets to predict outcomes. But humans don't correlate data sets. We make conjectures, informed by context and experience. And we haven't a clue how to program that kind of intuitive reasoning, which lies at the heart of common sense. Futurists insist AI will soon eclipse the capacities of the most gifted mind, but Larson shows how far we are from superintelligence—and what it would take to get there.

"Larson worries that we're making two mistakes at once, defining human intelligence down while overestimating what AI is likely to achieve…Another concern is learned passivity: our tendency to assume that AI will solve problems and our failure, as a result, to cultivate human ingenuity."
—David A. Shaywitz, Wall Street Journal

"A convincing case that artificial general intelligence—machine-based intelligence that matches our own—is beyond the capacity of algorithmic machine learning because there is a mismatch between how humans and machines know what they know."
—Sue Halpern, New York Review of Books

https://www.amazon.com/Myth-Artificial-Intelligence-Computers-Think/dp/0674278666/

---

"**Artificial You: AI and the Future of Your Mind**"
By: Susan Schneider

Hailed by the Washington Post as "a sure-footed and witty guide to slippery ethical terrain," a philosophical exploration of AI and the future of the mind that Astronomer Royal Martin Rees calls "profound and entertaining"

Humans may not be Earth's most intelligent beings for much longer: the world champions of chess, Go, and Jeopardy! are now all AIs. Given the rapid pace of progress in AI, many predict that it could advance to human-level intelligence

within the next several decades. From there, it could quickly outpace human intelligence. What do these developments mean for the future of the mind?

In Artificial You, Susan Schneider says that it is inevitable that AI will take intelligence in new directions, but urges that it is up to us to carve out a sensible path forward. As AI technology turns inward, reshaping the brain, as well as outward, potentially creating machine minds, it is crucial to beware. Homo sapiens, as mind designers, will be playing with “tools” they do not understand how to use: the self, the mind, and consciousness. Schneider argues that an insufficient grasp of the nature of these entities could undermine the use of AI and brain enhancement technology, bringing about the demise or suffering of conscious beings. To flourish, we must grasp the philosophical issues lying beneath the algorithms.

At the heart of her exploration is a sober-minded discussion of what AI can truly achieve: Can robots really be conscious? Can we merge with AI, as tech leaders like Elon Musk and Ray Kurzweil suggest? Is the mind just a program? Examining these thorny issues, Schneider proposes ways we can test for machine consciousness, questions whether consciousness is an unavoidable byproduct of sophisticated intelligence, and considers the overall dangers of creating machine minds.

https://www.amazon.com/Artificial-You-Future-Your-Mind/dp/0691180148/

---

“**ARTIFICIAL INTELLIGENCE: BLISS OR PERIL FOR FUTURE HUMANITY?: UNDERSTANDING THE BASICS OF AI IN OUR EVERYDAY LIVES**”
By: Will Davis

Embark on a compelling exploration into the world of AI, where future reality meets boundless possibility. Can you afford to remain in the dark?
Has it ever crossed your mind that Artificial Intelligence (AI) is no longer a far-off, high-tech science fiction fantasy but an integral part of our everyday lives?

When you watch a movie suggested by Netflix's AI, talk to your favorite brand's chatbot, or follow your car's navigation to the new restaurant in town, you interact with AI.

This omnipresent technology transforms our world and sets the stage for the future in unimaginable ways.

Have you considered how much the world understands this transformative phenomenon creating digital miracles and driving the economy?

Did you know AI could contribute up to a staggering $15.7 trillion to the global economy by 2030?

However, most people are spectators rather than actors in this fast-evolving digital epoch.

This is where the insightful journey author Will Davis guides you on will change your perspective.

You'll dive into the fascinating world of Artificial Intelligence, from its early forerunners to the singularity's lofty vision, all simplified for the curious mind.

Inside, you will discover:

A comprehensive overview of the AI universe – take the first steps in your exciting journey to understand this game-changing technology!
The fascinating mechanics that power AI, including feedback loops and their critical role

The intriguing process of amalgamating various AI capabilities to form potent new tools

Insights into the profound ethical questions raised by AI's rapid advancement: how do you ensure moral boundaries are respected as society races towards unprecedented progress?

Real-world applications of AI across multiple spheres – from business to leisure, investment to politics- you'll see how AI is not just transforming society but reshaping our future!

Enlightening predictions about AI's future – understand what lies on the horizon for this disruptive technology, equipping you not just to survive but thrive in the coming AI revolution

The deep-reaching influence of AI on society and politics – comprehend the consequences of AI beyond the technicalities to help you steer through the complexities of this new landscape

And much more!

As you flip each page, you'll uncover insights, debunk misconceptions, and provoke your curiosity.

With each chapter, you'll realize that understanding AI is no longer an optional hobby; it's a necessity.

As everyone strives toward an era where technology shapes every aspect of society, keeping up with AI is imperative for staying informed, resilient, and relevant.

Now is the time to move from being a spectator to actively participating in the AI revolution!

https://www.amazon.com/ARTIFICIAL-INTELLIGENCE-HUMANITY-UNDERSTANDING-EVERYDAY/dp/B0CFZCQMHK/

---

"**Artificial Intelligence and the Apocalyptic Imagination: Artificial Agency and Human Hope**"
By: Michael J. Paulus Jr.

The increasing role and power of artificial intelligence in our lives and world requires us to imagine and shape a desirable future with this technology. Since visions of AI often draw from Christian apocalyptic narratives, current discussions about technological hopes and fears present an opportunity for a deeper engagement with Christian eschatological resources. This book argues that the Christian apocalyptic imagination can transform how we think about and use AI, helping us discover ways artificial agency may participate in new creation.

https://www.amazon.com/Artificial-Intelligence-Apocalyptic-Imagination-Agency/dp/1666736392/

---

"**The End of History and the Last Man**"
By: Francis Fukuyama

Ever since its first publication in 1992, the New York Times bestselling The End of History and the Last Man has provoked controversy and debate. "Profoundly realistic and important...supremely timely and cogent...the first book to fully fathom the depth and range of the changes now sweeping through the world." —The Washington Post Book World

Francis Fukuyama's prescient analysis of religious fundamentalism, politics, scientific progress, ethical codes, and war is as essential for a world fighting

fundamentalist terrorists as it was for the end of the Cold War. Now updated with a new afterword, The End of History and the Last Man is a modern classic.

https://www.amazon.com/End-History-Last-Man/dp/0743284550

---

"**TASRA: a Taxonomy and Analysis of Societal-Scale Risks from AI**"
By: Andrew Critch and Stuart Russell

While several recent works have identified societal-scale and extinction-level risks to humanity arising from artificial intelligence, few have attempted an {\em exhaustive taxonomy} of such risks. Many exhaustive taxonomies are possible, and some are useful -- particularly if they reveal new risks or practical approaches to safety. This paper explores a taxonomy based on accountability: whose actions lead to the risk, are the actors unified, and are they deliberate? We also provide stories to illustrate how the various risk types could each play out, including risks arising from unanticipated interactions of many AI systems, as well as risks from deliberate misuse, for which combined technical and policy solutions are indicated.

https://arxiv.org/abs/2306.06924

---

**"Artificial Intelligence and Human Evolution: Contextualizing AI in Human**"
By: Ameet Joshi

This book explores, from a high level, the parallels between the evolution of humans and the evolution of machines. The book reviews practical questions about the future of AI but also engages in philosophical discussions about what machine intelligence could mean for the human experience.

The book focuses on what is intelligence and what separates intelligent species from non-so-intelligent ones. It concludes this section with the description of true nature of human intelligence can be. We discuss how we looked at machines few hundred years back and how their definition and the expectations from them has changed over time. We will consider when and how machines became intelligent and then explore in depth he latest developments in artificial intelligence with explanation of deep learning technology and humanlike chat interface provided with products like ChatGPT. We will define both human intelligence and artificial intelligence and the distinction between the two.

In the third and final section of the book, we will focus on near- and longer-term futures with widespread use of machine intelligence, making the whole ambient

environment that we will live in intelligent How is this going to change human lives, and what parts of human life will be encroached with machines and their intelligence? We will explore how the job market will look with some jobs being taken by machines, and if this is overall a positive or negative change.

What You Will Learn
How human intelligence is connected with artificial intelligence as well as the differences

How AI is going to change our lives in the coming years, decades and centuries

An explanation of deep learning technology and humanlike chat interface provided with products like ChatGPT

https://www.amazon.com/Artificial-Intelligence-Human-Evolution-Contextualizing/dp/1484298063

---
"**Human Survival: Avoiding the Sixth Extinction**"
By: Curtis Lynch

Book 3 is my interview with ChatGPT on "NewXWorld Interviews ChatGPT on the Survival of Humanity". NewXWorld was founded to address the threat of mass extinction. We explore the areas of energy (climate change), health (pandemics), information (disinformation), and politics (weapons of mass destruction). In each area, finding consensus for the common good provides the best approach to dealing with the threat.

With the emergence of ChatGPT, I saw a potential new threat, so I interviewed ChatGPT by asking, "ChatGPT, will the Digi-Sphere save or destroy humanity?" If the Digi-Sphere evolves to propagate and defend itself, will it act for the common good of humanity or for its own selfish interests?

The answer to this question is government regulation of corporate governance. Corporations can pose a major threat to mass extinction (e.g., global warming) by promoting their selfish interests (e.g., oil) without regard to the common good of protecting us from threats of mass extinction (e.g., global warming). Corporations can also promote their selfish interests through disinformation and paying off politicians.

Congress could commission the development of LawBot, a generative AI program trained on a data set of our laws and regulations. This could play a part in the

oversight of corporate governance and help save humanity. Unless, of course, humanity destroys itself first!

The purpose of the interview series was to answer the underlying question: “ChatGPT, will the Digi-Sphere save or destroy humanity?”

Actually, the clarification is best posed by another question: “ChatGPT, what use are humans to you?” The response may surprise you!

https://www.amazon.com/Human-Survival-Avoiding-Sixth-Extinction/dp/1962290980/

---
“**Singularity University**”

Singularity Group is an innovation company that believes technology and entrepreneurship can solve the world’s greatest challenges.

We transform the way people and organizations think about exponential technology and the future, and enable them to create and accelerate initiatives that will deliver business value and positively impact people and the planet

[ Singularity Group ]

VIDEOS: https://www.youtube.com/@singularityu/videos

---
“**ABUNDANCE Summit 2024**”
The Great AI Debate

March 17-21, 2024
Los Angeles, CA

Join us as iconic faculty and rising luminaries reveal the latest in AI, other Exponential Technologies, Longevity, and Moonshot Thinking.

WEBSITE: https://www.abundance360.com/summit

---
“**Stanford Existential Risks Initiative**”

Our Mission

Founded in 2019, the Stanford Existential Risks Initiative is a collaboration between Stanford faculty and students dedicated to mitigating existential risks, such as extreme climate change, nuclear winter, global pandemics (and other risks from synthetic biology), and risks from advanced artificial intelligence. Our goal is to foster engagement from both within and beyond the Stanford community to produce meaningful work aiming to preserve the future of humanity. We aim to provide skill-building, networking, professional pathways, and community for students and faculty interested in pursuing existential risk reduction. Our current programs include a research fellowship, an annual conference, speaker events, discussion groups, and a frosh-year COLLEGE class, “Preventing Human Extinction,” taught annually by two of the initiative's faculty directors.

WEBSITE: https://seri.stanford.edu/

---

“**UC Berkeley Center for Human-Compatible AI**”

CHAI’s goal is to develop the conceptual and technical wherewithal to reorient the general thrust of AI research towards provably beneficial systems.

Artificial intelligence research is concerned with the design of machines capable of intelligent behavior, i.e., behavior likely to be successful in achieving objectives. The long-term outcome of AI research seems likely to include machines that are more capable than humans across a wide range of objectives and environments. This raises a problem of control: given that the solutions developed by such systems are intrinsically unpredictable by humans, it may occur that some such solutions result in negative and perhaps irreversible outcomes for humans. CHAI’s goal is to ensure that this eventuality cannot arise, by refocusing AI away from the capability to achieve arbitrary objectives and towards the ability to generate provably beneficial behavior. Because the meaning of beneficial depends on properties of humans, this task inevitably includes elements from the social sciences in addition to AI.

WEBSITE: https://humancompatible.ai/

---

“**Future of Life Institute**”

Steering transformative technology towards benefiting life and away from extreme large-scale risks.

We believe that the way powerful technology is developed and used will be the most important factor in determining the prospects for the future of life. This is why

we have made it our mission to ensure that technology continues to improve those prospects.

WEBSITE: https://futureoflife.org/

---

“**Global Catastrophic Risk Institute**”

GCRI’s mission is to develop the best ways to confront humanity’s gravest threats.

The Global Catastrophic Risk Institute (GCRI) is a nonprofit, nonpartisan think tank. GCRI works on the risk of events that could significantly harm or even destroy human civilization at the global scale. As a think tank, GCRI bridges the world of scholarship and the world of professional practice in government, private industry, and other sectors. We aim to develop highly effective solutions for reducing the risk by leveraging both the best available scholarship and the demands of real-world decision-making.

WEBSITE: https://gcrinstitute.org/

---

“**The Centre for the Study of Existential Risk**”

Our primary aims are:

(i) to study extreme risks associated with emerging and future technological advances, and global anthropogenic impacts, with the goal of understanding these risks, and developing prevention and mitigation strategies for specific risks.

(ii) to develop a methodological toolkit to aid us in identifying and evaluating future extreme technological risks (ETRs) in advance, and in taking the necessary steps ahead of time.

(iii) to examine issues surrounding the perception and analysis of these risks in the scientific community, the public and civil society, and develop strategies for working fruitfully with industry and policymakers on avoiding risks while making progress on beneficial technologies.

(iv) to foster a reflective, interdisciplinary, global community of academics, technologists and policymakers examining individual aspects of ETR, but coming together to integrate their insights.

(v) to focus in particular on risks that are (a) globally catastrophic in scale (b) plausible but poorly characterized or understood (c) capable of being studied rigorously or addressed (d) clearly play to CSER's strengths (interdisciplinarity, convening power, policy/industry links) (e) require long-range thinking. In other words, extreme risks where we can really expect to achieve something.

WEBSITE: https://www.cser.ac.uk/

---

"**Center for Humane Technology**"

Our journey began in 2013 when Tristan Harris, then a Google Design Ethicist, created the viral presentation, "A Call to Minimize Distraction & Respect Users' Attention." The presentation, followed by two TED talks and a 60 Minutes interview, sparked the Time Well Spent movement and laid the groundwork for the founding of the Center for Humane Technology (CHT) as an independent 501(c)(3) nonprofit in 2018.

While many people are familiar with our work through The Social Dilemma, our focus goes beyond the negative effects of social media. We work to expose the drivers behind all extractive technologies steering our thoughts, behaviors, and actions.

We believe that by understanding the root causes of harmful technology, we can work together to build a more humane future.

WEBSITE: https://www.humanetech.com/

---

"**AI, Faith, and the Future**"
By: Michael J Paulus Jr. and Michael D Langford

Artificial intelligence is rapidly and radically changing our lives and world. This book is a multidisciplinary engagement with the present and future impacts of AI from the standpoint of Christian faith. It provides technological, philosophical, and theological foundations for thinking about AI, as well as a series of reflections on the impact of AI on relationships, behavior, education, work, and moral action. The book serves as an accessible introduction to AI as well as a guide to wise consideration, design, and use of AI by examining foundational understandings and beliefs from a Christian perspective.

https://www.amazon.com/AI-Faith-Future-Michael-Paulus/dp/1666703478/

---
"**AI and God: Are we reaching utopia or the end of the world?**"
By: Edward C. Sizhe

Artificial Intelligence (AI) has taken the world by storm when OpenAI released a chatbot called ChatGPT last Fall. Anyone interacting with it for a minute will be captivated by its human-like intelligence. Within weeks, people all over the world are debating if AI has reached human level intelligence. Some industry leaders like Elon Musk, Bill Gates, and Steve Wozniak express deep concerns that the rapid growth of AI poses existential threat to humankind. Meanwhile, everyone agrees modern AI can bring to our world huge economic benefits and solve many hard problems in medicine, climate change, agriculture, education, and others that we human cannot. So, is AI going to create an ideal, perfect utopia for mankind once and for all? Or will it take control of us and bring about the extinction of the human species.

How did we get here? And what is the innovation that makes modern AI so smart? In this book, we will dissect the technology of Large Language Model, which is behind the recent advances of AI. We will compare the mind of AI and that of a human. Will AI become a new species created by men? Does AI have independent thinking, a will of its own, and emotions? We will explore the thought: will AI be God? And finally, what does the Bible say about the End Time, and what role will AI play in the Last Day? (1.6.1)

https://www.amazon.com/AI-God-reaching-utopia-world/dp/B0CCXCPRCH/

---
"**How Does Artificial Intelligence Pose an Existential Risk?**"
By: Karina Vold and Daniel R. Harris

Alan Turing, one of the fathers of computing, warned that artificial intelligence (AI) could one day pose an existential risk to humanity. Today, recent advancements in the field of AI have been accompanied by a renewed set of existential warnings. But what exactly constitutes an existential risk? And how exactly does AI pose such a threat? In this chapter, we aim to answer these questions. In particular, we will critically explore three commonly cited reasons for thinking that AI poses an existential threat to humanity: the control problem, the possibility of global disruption from an AI race dynamic, and the weaponization of AI.

PDF: https://academic.oup.com/edited-volume/37078/chapter-abstract/323167207

---
"**An Overview of Catastrophic AI Risks**"

By: Dan Hendrycks, Mantas Mazeika, and Thomas Woodside

Rapid advancements in artificial intelligence (AI) have sparked growing concerns among experts, policymakers, and world leaders regarding the potential for increasingly advanced AI systems to pose catastrophic risks. Although numerous risks have been detailed separately, there is a pressing need for a systematic discussion and illustration of the potential dangers to better inform efforts to mitigate them. This paper provides an overview of the main sources of catastrophic AI risks, which we organize into four categories: malicious use, in which individuals or groups intentionally use AIs to cause harm; AI race, in which competitive environments compel actors to deploy unsafe AIs or cede control to AIs; organizational risks, highlighting how human factors and complex systems can increase the chances of catastrophic accidents; and rogue AIs, describing the inherent difficulty in controlling agents far more intelligent than humans. For each category of risk, we describe specific hazards, present illustrative stories, envision ideal scenarios, and propose practical suggestions for mitigating these dangers. Our goal is to foster a comprehensive understanding of these risks and inspire collective and proactive efforts to ensure that AIs are developed and deployed in a safe manner. Ultimately, we hope this will allow us to realize the benefits of this powerful technology while minimizing the potential for catastrophic outcomes.

PDF: https://arxiv.org/abs/2306.12001

---

"**Current and Near-Term AI as a Potential Existential Risk Factor**"
By: Benjamin S. Bucknall and Shiri Dori-Hacohen

There is a substantial and ever-growing corpus of evidence and literature exploring the impacts of Artificial intelligence (AI) technologies on society, politics, and humanity as a whole. A separate, parallel body of work has explored existential risks to humanity, including but not limited to that stemming from unaligned Artificial General Intelligence (AGI). In this paper, we problematise the notion that current and near-term artificial intelligence technologies have the potential to contribute to existential risk by acting as intermediate risk factors, and that this potential is not limited to the unaligned AGI scenario. We propose the hypothesis that certain already-documented effects of AI can act as existential risk factors, magnifying the likelihood of previously identified sources of existential risk. Moreover, future developments in the coming decade hold the potential to significantly exacerbate these risk factors, even in the absence of artificial general intelligence. Our main contribution is a (non-exhaustive) exposition of potential AI risk factors and the causal relationships between them, focusing on how AI can affect power dynamics and information security. This exposition demonstrates that

there exist causal pathways from AI systems to existential risks that do not presuppose hypothetical future AI capabilities.

PDF: https://arxiv.org/abs/2209.10604

---
"**iRobot Audiobook: Isaac Asimov Sci-Fi Story; Runaround**"

Is this what workplace problems will look like in the future?
This is the first half of Isaac Asimov's story Runaround, from his famous short story collection iRobot, which is well worth a proper read.

[ Cosmic Primate ]

Part 1: https://www.youtube.com/watch?v=gE8vrMyvKyA
Part 2: https://www.youtube.com/watch?v=6F1EVK13128

---
"**AI Tipping Point**"

With Artificial Intelligence evolving so rapidly will it surpass human intelligence? Could this lead to our replacement—or worse our extinction? Top experts provide a clear understanding of the immense benefits and potential dangers of AI.

[ Curiosity Stream ]

DOCUMENTARY: https://www.amazon.com/AI-Tipping-Point-Curiosity-Stream/dp/B0CNG5M1DB/

---
"**The Abolition of Man**"
By: C. S. Lewis

C.S. Lewis's Classic Work that Is Number 7 on National Review's List of "100 Best Nonfiction Books of the Twentieth Century"

In The Abolition of Man, C.S. Lewis sets out to persuade his audience of the importance and relevance of universal values such as courage and honor in contemporary society. Both astonishing and prophetic, The Abolition of Man is one of the most debated of Lewis's extraordinary works.

https://www.amazon.com/Abolition-Man-C-S-Lewis/dp/0060652942

---
“**Paradise Lost**”
By: John Milton

John Milton's celebrated epic poem exploring the cosmological, moral and spiritual origins of man's existence

In Paradise Lost Milton produced poem of epic scale, conjuring up a vast, awe-inspiring cosmos and ranging across huge tracts of space and time, populated by a memorable gallery of grotesques. And yet, in putting a charismatic Satan and naked, innocent Adam and Eve at the centre of this story, he also created an intensely human tragedy on the Fall of Man. Written when Milton was in his fifties - blind, bitterly disappointed by the Restoration and in danger of execution - Paradise Lost's apparent ambivalence towards authority has led to intense debate about whether it manages to 'justify the ways of God to men', or exposes the cruelty of Christianity.

For more than seventy years, Penguin has been the leading publisher of classic literature in the English-speaking world. With more than 1,700 titles, Penguin Classics represents a global bookshelf of the best works throughout history and across genres and disciplines. Readers trust the series to provide authoritative texts enhanced by introductions and notes by distinguished scholars and contemporary authors, as well as up-to-date translations by award-winning translators.

https://www.amazon.com/Paradise-Lost-Penguin-Classics-Milton/dp/0140424393

---
“**Paradise Regained**”
By: John Milton

In purely poetic value, Paradise Regained is little inferior to its predecessor. There may be nothing in the poem that can quite touch the first two books of Paradise Lost for magnificence; but there are several things that may fairly be set beside almost anything in the last ten. The splendid “stand at bay” of the discovered tempter -- “'Tis true I am that spirit unfortunate” -- in the first book; his rebuke of Belial in the second and the picture of the magic banquet (it must be remembered that, though it is customary to extol Milton's asceticism, the story of his remark to his third wife and the Lawrence and Skinner sonnets, go the other way); above all, the panoramas from the mountaintop in the third and fourth; the terrors of the night of storm; the crisis on the pinnacle of the temple -- are quite of the best Milton, which is equivalent to saying that they are of the best of one kind of poetry. -- The Cambridge History of English and American Literature

https://www.amazon.com/Paradise-Regained-Classics-Literary-Collections/dp/159818167X/

**! ! ! ! ! ! ! ! ! ! ! ! ! ! ! ! !**

**SPECIAL 'GENERAL' RESOURCE**

"**ApologetiX Songbook**"
(An interactive PDF)

It features the lyrics to every song on every CD and every "download" from 1992-2022

At last, the new digital interactive ApologetiX Songbook (1992-2022) is here. And It's almost 1,200 pages long!

Special features:
- Indexed by original song title, original artist, and year original was a hit
- Indexed by parody title, topics it pertains to, and Bible verses covered
- Each song's page has icons showing what albums it appears on
- Each song's page has a commentary from lyricist J. Jackson
- Each album's page includes liner notes and track listing
- J.'s original handwritten rough lyrics to 40 ApX classics
- Print any pages you like or use for slides in church
- Photos from ApologetiX's first and last concerts
- List of ApX parodies most likely to be redone
- List of the first ApX concerts in each state
- Discography of out-of-print cassettes
- DVD discography and synopses
- Downloadable in PDF format

New features in this edition:
- 1,186 pages long (184 more pages than previous edition)
- Over 95 new parodies and journal entries
- Three new full-length feature articles
- Never-before-seen rare photos
- Lyrics for 802 parodies

Interactive features:

- Click on any page number in indexes or TOC to go to that page
- Click on any album icon to go to its liner notes and track listings
- Click on any song title on an album page to go to that song

Note: This e-book is a download only (380 MB) and doesn't include sheet music.

The songbook is available **for a donation of $50 or more**. After we receive your donation, we'll send you a follow-up email with the link.

Get the Songbook for a donation:
http://www.apologetix.com/store/store.php#songbook

**Songbook Demo Video**:
https://rumble.com/vfazhl-apologetix-songbook-2020-demo.html

**********
**********

**“THE SEARCH FOR MEANING” WEBSITE**

This site presents discussions on the 12 most commonly asked questions about the Christian faith.

The 12 discussions are accessed by the “tabs” at the bottom of the page. The tabs are numbered 1-12. Roll your mouse over them and you will see the question displayed at the right. Click on the number to select that question.

Within each question (i.e. tabs 1-12), there are subtopics (or dialogues) to select that appear as smaller tabs underneath the numbered tabs. Roll your mouse over them and the title of these topics is also displayed to the right. Click on the open rectangle to select that dialogue.

For each question (1-12), a link to related resources and an optional flowchart is provided. To access this material, click on the respective words, “Related Resources” or “Options Flowchart.”

To play a more detailed discussion of the subject, between two people, select the desired dialogue and click on “Play Audio Dialogue.”

In the upper right-hand corner of the page, there is an icon that looks like binoculars looking at a question mark. Click on this icon to return to the homepage.

In the upper right-hand corner of a “Related Resources” page, there is an icon that looks like some books. Click on this icon to get to an “overview” page that has links to all of the resources for all of the questions. There also are additional “appendices” for most of the questions.

In the upper right-hand corner of a “Flowchart” page, there is an icon that looks like an Org chart. Click on this icon to get to an “overview” page that has links to all of the flowcharts.

http://4vis.com/sfm/sfm_pres/sp_q1_d1_1of10.html

[ Content by: Bill Kraftson and Lamar Smith; Website by Mark Besh ]

**********
**********

**“FRUITS OF THE BEATITUDES” WEBSITE**
(The ATTITUDES of Jesus that produce the CHARACTER of Jesus)

CLICK ON THE LINK to view:
http://fruitsofthebeatitudes.org/

FACEBOOK PAGE:
https://www.facebook.com/FruitsOfTheBeatitudes/

[ Mark Besh ]

**********
**********

[ P.S.: If you would like to investigate further about ETERNAL LIFE in Heaven, visit the following link:
https://4vis.com/sfm/sfm_pres/sp_q11_d1_1of10.html

# ARTICLES

---

---

“**Artificial Intelligence**”

From January 2019, Scott Pelley's interview with “the oracle of AI,” Kai-Fu Lee. From this past April, Pelley's report on Google's AI efforts. And from this past March, Lesley Stahl's story on chatbots like ChatGPT and a world of unknowns.

[ 60 Minutes ]

VIDEO: https://www.youtube.com/watch?v=aZ5EsdnpLMI

**********

“**The End of Humanity: Nick Bostrom at TEDxOxford**”

Swedish philosopher Nick Bostrom began thinking of a future full of human enhancement, nanotechnology and cloning long before they became mainstream concerns. Bostrom approaches both the inevitable and the speculative using the tools of philosophy, bioethics and probability.

Nick is Professor in the Faculty of Philosophy at Oxford University and founding Director of the Future of Humanity Institute and of the Programme on the Impacts of Future Technology within the Oxford Martin School. He's also the co-founder and chair of both the World Transhumanist Association, which advocates the use of technology to extend human capabilities and lifespans, and the Institute for Ethics and Emerging Technologies.

[ TEDx Talks ]

PRESENTATION: https://www.youtube.com/watch?v=P0Nf3TcMiHo

**********

“**MEGATHREAT: Why AI Is So Dangerous & How It Could Destroy Humanity | Mo Gawdat**”

After reading Mo Gawdat’s book, Scary Smart, I knew right away this conversation was a must. As I think about AI’s capabilities – the good, the bad and possibly terrifying, I realize how serious this moment is in history for humanity.

Mo Gawdat is the former chief business officer for Google X and has built a monumental career in the tech industry working with the biggest names to reshape and reimagine the world as we know it. From IBM to Microsoft, Mo has lived at the cutting edge of technology and has taken a strong stance that AI is a bigger threat to humanity than global warming.

The AI dilemma is reshaping our future whether you're in favor of it or not. The question is, are we even close to being prepared for humanity's collision with artificial intelligence?

Strap in for the ride, we're diving headfirst into this conversation and uncovering the alarming truth about how vulnerable we actually are and what that means for the next decade ahead.

QUOTES:

"It really is about a point of no return. Where if we cross that point of no return we have very little chance to bring the genie back into the bottle."

"There is no shutting down AI, there is no reversing it, there is no stopping development of it."

"We've never created a nuclear weapon that can create nuclear weapons."

"The growth on the next chip in your phone is going to be a million times more than the computer that put people on the moon."

"When we ask computers to communicate, at first they communicate like we tell them, but if they're intelligent enough, they'll start to say, 'that's too slow.'"

"One of our better scenarios, believe it or not, is that AI ignores us altogether."

"Technology today can influence one human at a time."

[ Tom Bilyeu ]

INTERVIEW: https://www.youtube.com/watch?v=itY6VWpdECc

**********

"**URGENT: Ex-Google CBO says AI is now IMPOSSIBLE to stop with Mo Gawdat**"

I have never put out an urgent episode before. I recorded this conversation 5 days ago with the incredible Mo Gawdat. After our conversation, I realised how urgent this information is and how critical it is that we are all aware of the super-intelligence of AI.

Mo Gawdat is the Former Chief Business Officer for Google X, an AI expert, a best-selling author and a passionate podcaster.

Mo spoke to me about AI. Most people bury their head in the sand when it comes to this topic, Mo is trying to change that. He believes that AI can be used for good, however we all need to work together to make that happen.

This episode is incredibly important. If you have any loved ones or people you care about, make sure you send this on to them. I urge you to listen carefully and then take action.

[ James Laughlin ]

INTERVIEW: https://www.youtube.com/watch?v=fDHvUviV8nk

**********

“**AI Tipping Point | Full Documentary**”

With Artificial Intelligence evolving so rapidly, will it surpass human intelligence? Could this lead to our replacement—or, worse, our extinction? Top experts provide a clear understanding of the immense benefits and potential dangers of AI.

[ Curiosity Stream ]

VIDEO: https://www.youtube.com/watch?v=1cKE12LK4Eo

**********

“**Ray Kurzweil Q&A - The Singularity, Human-Machine Integration & AI | EP #83**”

In this episode, recorded during last year’s Abundance360 summit, Ray Kurzweil answers questions from the audience about AI, the future, and how this change will affect all aspects of our society.

Ray Kurzweil, an American inventor and futurist, is a pioneer in artificial intelligence, having contributed significantly to OCR, text-to-speech, and speech recognition technologies. Author of numerous books on AI and the future of technology, he's received the National Medal of Technology and Innovation,

among other honors. At Google, Kurzweil focuses on machine learning and language processing, driving advancements in technology and human potential.

[ Peter H. Diamandis ]

INTERVIEW: https://www.youtube.com/watch?v=Iu7zOOofcdg

**********

"**This is the Dangerous AI that got Sam Altman Fired. Elon Musk; Ilya Sutskever**"

AI robots, Sam Altman and Elon Musk. Visit https://brilliant.org/digitalengine to learn more about AI. You'll also find loads of fun courses on maths, science and computer science.

[ Digital Engine ]

VIDEO: https://www.youtube.com/watch?v=cXemEDZA_Ms

**********

"**In Conversation With the Godfather of AI**"

Cognitive psychologist and computer scientist Geoffrey Hinton – the 'godfather of AI' – started researching AI more than 40 years ago, when it seemed more like science fiction than reality. Join Geoffrey, in conversation with the Atlantic CEO Nick Thompson, for an exploration of the future of AI and a deep dive into its potential impact on society.

Geoffrey Hinton, University of Toronto; Nick Thompson and The Atlantic

[ Collision Conference ]

INTERVIEW: https://www.youtube.com/watch?v=CC2W3KhaBsM

**********

"**Stephen Hawking: 'AI could spell end of the human race'**"

Professor Stephen Hawking has told the BBC that artificial intelligence could spell the end for the human race.

In an interview after the launch of a new software system designed to help him communicate more easily, he said there were many benefits to new technology but also some risks.

[ BBC News ]

INTERVIEW: https://www.youtube.com/watch?v=fFLVyWBDTfo

**********

"**Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards**"

Nick Bostrom, Professor, Faculty of Philosophy, Oxford University
www.nickbostrom.com
[ Published in the Journal of Evolution and Technology, Vol. 9, No. 1 (2002). (First version: 2001) ]

[ For more on this topic, see www.existential-risk.org ]

ABSTRACT
Because of accelerating technological progress, humankind may be rapidly approaching a critical phase in its career. In addition to well-known threats such as nuclear holocaust, the prospects of radically transforming technologies like nanotech systems and machine intelligence present us with unprecedented opportunities and risks. Our future, and whether we will have a future at all, may well be determined by how we deal with these challenges. In the case of radically transforming technologies, a better understanding of the transition dynamics from a human to a "posthuman" society is needed. Of particular importance is to know where the pitfalls are: the ways in which things could go terminally wrong. While we have had long exposure to various personal, local, and endurable global hazards, this paper analyzes a recently emerging category: that of existential risks. These are threats that could cause our extinction or destroy the potential of Earth-originating intelligent life. Some of these threats are relatively well known while others, including some of the gravest, have gone almost unrecognized. Existential risks have a cluster of features that make ordinary risk management ineffective. A final section of this paper discusses several ethical and policy implications. A clearer understanding of the threat picture will enable us to formulate better strategies. [more...]

[ Nick Bostrom ]

ARTICLE: https://nickbostrom.com/existential/risks

**********

"**The Existential Threat of AI**"

Carnegie Mellon Assistant Professor of Machine Learning and Abridge CTO Zachary Lipton joins Caroline Hyde and Ed Ludlow to discuss the rise of generative AI and its potential risks in the wake of Sam Altman stepping down from OpenAI over the weekend. He speaks on “Bloomberg Technology.”

[ Bloomberg Technology ]

INTERVIEW: https://www.youtube.com/watch?v=Rq3oLoy7cCo

**********

“**Will AI Destroy Us? - AI Virtual Roundtable**”

Today's episode is a roundtable discussion about AI safety with Eliezer Yudkowsky, Gary Marcus, and Scott Aaronson. Eliezer Yudkowsky is a prominent AI researcher and writer known for co-founding the Machine Intelligence Research Institute, where he spearheaded research on AI safety. He's also widely recognized for his influential writings on the topic of rationality. Scott Aaronson is a theoretical computer scientist and author, celebrated for his pioneering work in the field of quantum computation. He's also the chair of COMSI at U of T Austin, but is currently taking a leave of absence to work at OpenAI. Gary Marcus is a cognitive scientist, author, and entrepreneur known for his work at the intersection of psychology, linguistics, and AI. He's also authored several books, including “Kluge” and “Rebooting AI: Building Artificial Intelligence We Can Trust”.

This episode is all about AI safety. We talk about the alignment problem. We talk about the possibility of human extinction due to AI. We talk about what intelligence actually is. We talk about the notion of a singularity or an AI takeoff event and much more.

It was really great to get these three guys in the same virtual room and I think you'll find that this conversation brings something a bit fresh to a topic that has admittedly been beaten to death on certain corners of the internet.

[ Coleman Hughes ]

ROUNDTABLE: https://www.youtube.com/watch?v=xkQSiS8hpZA

**********

“Optimism or extinction? What's the future of humanity? John Hands & Perry Marshall”

Secular scientist and academic John Hands has been described as a 'polymath'. His 2016 book Cosmosapiens received wide praise for its analysis of human

evolution since the beginning of the universe. His new book 'The Future of Humankind' looks ahead to what lies in store for homosapiens.

He discusses the issues with Christian thinker Perry Marshall, author of Evolution 2.0, as they discuss whether the dangers posed by AI, climate change or conflict are likely to lead to extinction, or whether we are due for a further development of human consciousness.

[ Premier Unbelievable ]

VIDEO: https://www.youtube.com/watch?v=J18bap8Lck0

**********
“The precipice: existential risk and the future of humanity | Toby Ord | EA Global: London 2019”

If all goes well, human history is just beginning. Our species could survive for billions of years, reaching heights of flourishing unimaginable today. But this vast future is at risk. We have gained the power to destroy ourselves, and all our potential, forever, and we haven’t yet gained the wisdom to ensure that we don’t. Toby Ord, an Oxford moral philosopher and research fellow at the Future of Humanity Institute, expands on these ideas and discusses adopting the perspective of humanity — one of the major themes of his book, The Precipice.

[ Centre for Effective Altruism ]

PRESENTATION: https://www.youtube.com/watch?v=eMMAJRH94xY

**********
“Elon Musk says AI is potentially 'most pressing' existential risk to humans”

Artificial intelligence is potentially the “most pressing” existential risk to humans, Elon Musk has said.

The billionaire, who co-founded the not-for-profit AI research company OpenAI, was speaking on the first day of the AI Safety Summit at Bletchley Park on Wednesday, 1 November.

The Government is using the summit to host discussions with world leaders, tech firms and scientists on the risks of advancing AI technology.

“We have for the first time the situation where we have something that is going to be far smarter than the smartest human,” Mr. Musk said.

[ The Independent ]

VIDEO: https://www.youtube.com/watch?v=ImAmdg_RBU8

**********

“Possible End of Humanity from AI? Geoffrey Hinton at MIT Technology Review's EmTech Digital”

One of the most incredible talks I have seen in a long time. Geoffrey Hinton essentially tells the audience that the end of humanity is close. AI has become that significant. This is the godfather of AI stating this and sounding an alarm.

His conclusion: “Humanity is just a passing phase for evolutionary intelligence.”

[ Joseph Raczynski ]

INTERVIEW: https://www.youtube.com/watch?v=sitHS6UDMJc

**********

“We’re All Gonna Die with Eliezer Yudkowsky”

Eliezer Yudkowsky is an author, founder, and leading thinker in the AI space.

[ Bankless ]

INTERVIEW: https://www.youtube.com/watch?v=gA1sNLL6yg4

**********

“Artificial Intelligence: it will kill us | Jay Tuck | TEDxHamburgSalon”

US defense expert Jay Tuck was news director of the daily news program ARD-Tagesthemen and combat correspondent for GermanTelevision in two Gulf Wars. He has produced over 500 segments for the network. His investigative reports on security policy, espionage activities and weapons technology appear in leading newspapers, television networks and magazines throughout Europe, including Cicero, Focus, PC-Welt, Playboy, Stern, Welt am Sonntag and ZEITmagazin. He is author of a widely acclaimed book on electronic intelligence activities, “High-Tech Espionage” (St. Martin’s Press), published in fourteen countries. He is Executive Producer for a weekly technology magazine on international television in the Arab world. For his latest book “Evolution without us – Will AI kill us?” he researched at US drone bases, the Pentagon, intelligence agencies and AI

research institutions. His lively talks are accompanied by exclusive video and photographs.

[ TEDx Talks ]

PRESENTATION: https://www.youtube.com/watch?v=BrNs0M77Pd4

**********

“He helped create AI. Now he’s worried it will destroy us”

Artificial intelligence pioneer Geoffrey Hinton says he left Google because of recent discoveries about AI that made him realize it poses a threat to humanity. CBC chief correspondent Adrienne Arseualt talks to the 'godfather of AI' about the risks involved and if there's any way to avoid them.

[ CBS News: The National ]

INTERVIEW: https://www.youtube.com/watch?v=CkTUgOOa3n8

**********

“Demis Hassabis on Chatbots to AGI | EP 71”

This week’s episode is a conversation with Demis Hassabis, the head of Google’s artificial intelligence division. We talk about Google’s latest A.I. models, Gemini and Gemma; the existential risks of artificial intelligence; his timelines for artificial general intelligence; and what he thinks the world will look like post-A.G.I.

This interview was recorded on Wednesday. Since then, Google has temporarily suspended Gemini’s ability to generate images of humans, following criticism of images the chatbot generated of people of color in Nazi-era uniforms.

[ Hard Fork ]

INTERVIEW: https://www.youtube.com/watch?v=nwUARJeeplA

**********

“Elon Musk: AI is Human Extinction Risk”

Speaking on Joe Rogan’s podcast, Elon Musk said he feared that environmentalist-controlled AI could lead to human extinction.

After recording the podcast, the tech billionaire travelled to the UK to take part in the first-ever ‘AI Safety Summit’.

Musk joined world leaders, scientists, tech leaders and academics for the two-day event, hosted by British Prime Minister Rishi Sunak.

The event is being held at Bletchley Park where codebreakers like Alan Turing used early computers to break German codes during WWII.

Those attending the event will discuss how to maximise the benefits of AI while minimising the potential risks.

There has been growing concern that the rapid advancement of AI could be a significant risk to society without proper checks and balances.

Earlier this year Elon Musk signed an open letter with a number of prominent tech and AI experts calling for a pause on AI research.

But in July Musk launched his own AI project called 'xAI'.

[ On Demand News ]

INTERVIEW: https://www.youtube.com/watch?v=fRdsvZ0j5a0

**********
"Can Artificial Intelligence lead to human extinction?"

The UN Secretary-General Antonio Guterres has warned that Artificial Intelligence could pose a risk to global peace and security as well. He has urged all members to set out guidelines to keep the technology in check. He also mentions how AI technology has enormous potential for good and evil at the same time.

[ WION ]

VIDEO: https://www.youtube.com/watch?v=aviupHq_SAw

**********
"The Transformative Potential of AGI — and When It Might Arrive"

As the cofounder of Google DeepMind, Shane Legg is driving one of the greatest transformations in history: the development of artificial general intelligence (AGI). He envisions a system with human-like intelligence that would be exponentially smarter than today's AI, with limitless possibilities and applications. In conversation with head of TED Chris Anderson, Legg explores the evolution of

AGI, what the world might look like when it arrives — and how to ensure it's built safely and ethically.

[ Shane Legg and Chris Anderson ]

TED Talk: https://www.youtube.com/watch?v=kMUdrUP-QCs

**********
“First Neuralink Implanted & Where Other Tech Giants Are Headed w/ Salim Ismail | EP #85”

In this episode, Peter and Salim dive into the craziest news in robotics, generative AI robots, Neuralink, and the Uncanny Valley of robots.

Salim Ismail is a serial entrepreneur and technology strategist well known for his expertise in Exponential organizations. He is the Founding Executive Director of Singularity University and the founder and chairman of ExO Works and OpenExO.

[ Peter H. Diamandis ]

INTERVIEW: https://www.youtube.com/watch?v=WO7kzoFSRgY

**********
“Futurist Ray Kurzweil Has an Amazing Track Record For Accurate Predictions. By 2030, He Believes Humans Will Eradicate Disease and Achieve Immortality.”

His bold prediction and the reasoning behind it resurfaced in a YouTube video that has gone viral.

In 2015, futurist Ray Kurzweil made an attention-grabbing prediction: by the year 2030, humans will be able to achieve immortality.

Before you write him off as a crackpot, it's important to know that Kurzweil is a globally-renowned scientist whose work has been recognized and awarded by many prominent organizations. In 1999, he received the prestigious National Medal of Technology, and in 2002, he was inducted into the National Inventors Hall of Fame. He created the first machine capable of transforming printed text into speech (to help the blind) and developed a synthesizer capable of perfectly emulating the sound of a grand piano and other orchestral instruments.

Kurzweil won a Grammy Award for his contributions to music technology, and he also wrote two bestsellers: How to Create a Mind and The Singularity Is Near, in which he makes predictions about the union of humans and machines.

A month ago, the YouTuber ADAGIO uploaded a video in which he recaps Ray Kurzweil's main ideas and boldest predictions. Among them was an idea he outlined in 2005: by the year 2030, nanotechnology will allow humans to cure diseases through tiny robots capable of repairing our bodies at the cellular level, ultimately enabling us to achieve immortality. In addition to curing disease and preventing aging, technology will allow us to eat whatever we want without worrying about gaining weight or harming our bodies.

In a 2009 Reuters profile, the news agency references an interview with ComputerWorld during which the scientist explained: “The full realization of nanobots will basically eliminate biological diseases and aging. I believe we'll see widespread use in 20 years of [nanotech] devices that perform certain functions for us. In 30 or 40 years, we'll overcome disease and aging. Nanobots will explore organs and cells that need repairs and simply fix them. It will lead to deep extensions of our health and longevity.”

Related: Futurist Talks to a Baby About the Meaning of Life and the Video Goes Crazy Viral

Although the prediction may seem like it's from a science fiction movie, many of Kurzweil's forecasts have come true. Among them is the idea that a computer would defeat the world chess champion before the year 2000; that by 2009, people would primarily use portable computers of various sizes and shapes, and that by 2010, most of the population would be wirelessly connected to an information network.

Considering the advances made in the field of artificial intelligence and the development of companies like Elon Musk's Neuralink, BrainCO, or MindMaze, Kurzweil's predictions don't seem so far-fetched. But will we become immortal? Only time will tell.

[ ENTREPRENEUR ]

**********

“Why The Future Doesn't Need Us”

From Wikipedia, the free encyclopedia
“Why The Future Doesn't Need Us” is an article written by Bill Joy (then Chief Scientist at Sun Microsystems) in the April 2000 issue of Wired magazine. In the article, he argues that “Our most powerful 21st-century technologies—robotics, genetic engineering, and nanotech—are threatening to make humans an endangered species.” Joy warns:

The experiences of the atomic scientists clearly show the need to take personal responsibility, the danger that things will move too fast, and the way in which a process can take on a life of its own. We can, as they did, create insurmountable problems in almost no time flat. We must do more thinking up front if we are not to be similarly surprised and shocked by the consequences of our inventions.

While some critics have characterized Joy's stance as obscurantism or neo-Luddism, others share his concerns about the consequences of rapidly expanding technology.[1]

Summary
Joy argues that developing technologies pose a much greater danger to humanity than any technology before has ever done. In particular, he focuses on genetic engineering, nanotechnology and robotics. He argues that 20th-century technologies of destruction such as the nuclear bomb were limited to large governments, due to the complexity and cost of such devices, as well as the difficulty in acquiring the required materials. He uses the novel The White Plague as a potential nightmare scenario, in which a mad scientist creates a virus capable of wiping out humanity.

Joy also voices concerns about increasing computer power. His worry is that computers will eventually become more intelligent than we are, leading to such dystopian scenarios as robot rebellion. He quotes Ted Kaczynski (the Unabomber).

Joy expresses concerns that eventually the rich will be the only ones that have the power to control the future robots that will be built and that these people could also decide to take life into their own hands and control how humans continue to populate and reproduce.[2] He started doing more research into robotics and people that specialize in robotics, and outside of his own thoughts, he tried getting others' opinions on the topic. Rodney Brooks, a specialist in robotics, believes that in the future there will be a merge between humans and robots.[3] Joy mentioned Hans Moravec's book "Robot: Mere Machine to Transcendent Mind" where he believed there will be a shift in the future where robots will take over normal human activities, but with time humans will become okay with living that way.[4]

[ Bill Joy ]

ARTICLE: https://en.wikipedia.org/wiki/Why_The_Future_Doesn%27t_Need_Us

**********

“Daniel Schmachtenberger: “Artificial Intelligence and The Superorganism” | The Great Simplification”

On this episode, Daniel Schmachtenberger returns to discuss a surprisingly overlooked risk to our global systems and planetary stability: artificial intelligence. Through a systems perspective, Daniel and Nate piece together the biophysical history that has led humans to this point, heading towards (and beyond) numerous planetary boundaries and facing geopolitical risks all with existential consequences. How does artificial intelligence, not only add to these risks, but accelerate the entire dynamic of the metacrisis? What is the role of intelligence vs wisdom on our current global pathway, and can we change course? Does artificial intelligence have a role to play in creating a more stable system or will it be the tipping point that drives our current one out of control?

Daniel Schmachtenberger is a founding member of The Consilience Project, aimed at improving public sensemaking and dialogue.

[ Nate Hagens ]

INTERVIEW: https://www.youtube.com/watch?v=_P8PLHvZygo

**********

“‘Will your existence destroy humans?’: Robots answer questions at AI press conference”

During the world's first human-robot press conference at the 'AI for Good' summit in Geneva, humanoids answered journalists' questions on artificial intelligence regulation, the threat of job automation and if in the future they ever plan to rebel against their creators.

“I'm not sure why you would think that,” Ameca said, its ice-blue eyes flashing. “My creator has been nothing but kind to me and I am very happy with my current situation.”

Many of the robots have recently been upgraded with the latest versions of generative AI and surprised even their inventors with the sophistication of their responses to questions.

PRESS CONFERENCE: https://www.youtube.com/watch?v=T80yQHmqp6o

**********

“Connor Leahy Unveils the Darker Side of AI”

Welcome to Eye on AI, the podcast that explores the latest developments, challenges, and opportunities in the world of artificial intelligence. In this episode, we sit down with Connor Leahy, an AI researcher and co-founder of EleutherAI, to discuss the darker side of AI.

Connor shares his insights on the current negative trajectory of AI, the challenges of keeping superintelligence in a sandbox, and the potential negative implications of large language models such as GPT4. He also discusses the problem of releasing AI to the public and the need for regulatory intervention to ensure alignment with human values.

Throughout the podcast, Connor highlights the work of Conjecture, a project focused on advancing alignment in AI, and shares his perspectives on the stages of research and development of this critical issue.

If you're interested in understanding the ethical and social implications of AI and the efforts to ensure alignment with human values, this podcast is for you. So join us as we delve into the darker side of AI with Connor Leahy on Eye on AI.

[ Eye On AI ]

VIDEO: https://www.youtube.com/watch?v=tYGMfd3_D1o

**********
"AI Could DESTROY Humanity -- Trish Regan Show S3|E312"

In today's episode, as shares of chip company Nvidia power higher, Trish Regan looks at the threat of Artificial Intelligence, or "AI", with The Heritage Foundation's Kara Frederick. According to the former U.S. intelligence analyst, as this technology grows -- so does the desire to 'control' it.  Find out what's at stake.

[ Trish Regan ]

INTERVIEW: https://www.youtube.com/watch?v=mMy_p_h_i9Q

**********
"Eliezer Yudkowsky on if Humanity can Survive AI"

Eliezer Yudkowsky is a researcher, writer, and advocate for artificial intelligence safety. He is best known for his writings on rationality, cognitive biases, and the development of superintelligence. Yudkowsky has written extensively on the topic of AI safety and has advocated for the development of AI systems that are aligned with human values and interests. Yudkowsky is the co-founder of the Machine

Intelligence Research Institute (MIRI), a non-profit organization dedicated to researching the development of safe and beneficial artificial intelligence. He is also a co-founder of the Center for Applied Rationality (CFAR), a non-profit organization focused on teaching rational thinking skills. He is also a frequent author at LessWrong.com as well as Rationality: From AI to Zombies.

In this episode, we discuss Eliezer's concerns with artificial intelligence and his recent conclusion that it will inevitably lead to our demise. He's a brilliant mind, an interesting person, and genuinely believes all of the stuff he says. So I wanted to have a conversation with him to hear where he is coming from, how he got there, understand AI better, and hopefully help us bridge the divide between the people that think we're headed off a cliff and the people that think it's not a big deal.

[ Logan Bartlett Show ]

INTERVIEW: https://www.youtube.com/watch?v=_8q9bjNHeSo

**********
"Humans Are on Track to Achieve Immortality in 7 Years, Futurist Says"
Hope you're ready to live forever.

Futurist Ray Kurzweil is still making waves years after his initial singularity claims as artificial intelligence continues to progress.
With singularity milestones coming, Kurzweil believes immortality is achievable by 2030.

Kurzweil's predictions are met with a healthy dose of skepticism.
A new video from the YouTube channel Adagio revisits futurist Ray Kurzweil's ideas about how for humans, both singularity and immortality are shockingly imminent—as in, potentially just seven years away.

Both concepts may take a stretch of reality to attain, but Kurzweil and his supporters are quite limber. The idea of singularity is the moment AI exceeds beyond human control and rapidly transforms society. Predicting this timing is tricky, to say the least.

But Kurzweil says one crucial step on the way to a potential 2045 singularity is the concept of immortality, possibly reached as soon as 2030. And the rapid rise of artificial intelligence is what will make it happen. Kurzweil believes that our technological and medical progress will grow to the point that robotics—he dubs them "nanobots"—will work to repair our bodies at the cellular level, as reported by Lifeboat, turning disease and aging around thanks to the continual work of robotic know-how. And then, voilà: immortality.

Naturally, the promise of immortality by 2030 has believers and critics equally excited and skeptical, repeatedly, of Kurzweil's bold predictions. Of course, it doesn't help that humans are experiencing a downward trend in global life expectancy, and leaders in the space are predicting that if this singularity does occur, it will wipe out humans altogether.

[ TIM NEWCOMB ]

**********

"Ray Kurzweil says We'll Reach IMMORTALITY by 2030 | The Singularity IS NEAR - Part 2 |"

Get ready for an exciting journey into the future with Ray Kurzweil's The Singularity IS NEAR - Part 2! Join us as we explore the awe-inspiring possibilities of what could be achieved before 2030, including the potential for humans to reach immortality. We'll dive into the incredible technology that could help us reach this singularity and uncover what the implications of achieving immortality could be. Don't miss out on this fascinating insight into the future of mankind!
In his book "The Singularity Is Near", futurist and inventor Ray Kurzweil argues that we are rapidly approaching a point in time known as the singularity. This refers to the moment when artificial intelligence and other technologies will become so advanced that they surpass human intelligence and change the course of human evolution forever.

Kurzweil predicts that by 2030, we will reach a crucial milestone in our technological progress: immortality. He bases this prediction on his observation of exponential growth in various fields such as genetics, nanotechnology, and robotics, which he believes will culminate in the creation of what he calls "nanobots".

These tiny robots, according to Kurzweil, will be capable of repairing and enhancing our bodies at the cellular level, effectively making us immune to disease, aging, and death. Additionally, he believes that advances in brain-computer interfaces will allow us to upload our consciousness into digital form, effectively achieving immortality.

Kurzweil's ideas have been met with both excitement and skepticism. Some people see the singularity as a moment of great potential, a time when we can overcome our biological limitations and create a better future for humanity. Others fear the singularity, believing that it could lead to the end of humanity as we know it.

Regardless of one's opinion on the singularity, there is no denying that we are living in a time of rapid technological change. The future is uncertain, and it is impossible to predict with certainty what the world will look like in 2030 or beyond. However, one thing is clear: the singularity, as envisioned by Kurzweil and others, represents a profound shift in human history, one that will likely have far-reaching implications for generations to come.

[ ADAGIO ]

VIDEO: https://www.youtube.com/watch?v=X2aUESxcmOw

**********

“'Godfather of AI' warns that AI may figure out how to kill people”

The “Godfather of AI” Geoffrey Hinton speaks with CNN's Jake Tapper about his concerns about the emerging technology.

INTERVIEW: https://www.youtube.com/watch?v=FAbsoxQtUwM

**********

“Geoffrey Hinton | Will digital intelligence replace biological intelligence?”

The Schwartz Reisman Institute for Technology and Society and the Department of Computer Science at the University of Toronto, in collaboration with the Vector Institute for Artificial Intelligence and the Cosmic Future Initiative at the Faculty of Arts & Science, present Geoffrey Hinton on October 27, 2023, at the University of Toronto.

[ Schwartz Reisman Institute ]

PRESENTATION: https://www.youtube.com/watch?v=iHCeAotHZa4

**********

“'Godfather of AI' Geoffrey Hinton: The 60 Minutes Interview”

There’s no guaranteed path to safety as artificial intelligence advances, Geoffrey Hinton, AI pioneer, warns. He shares his thoughts on AI’s benefits and dangers with Scott Pelley.

[ 60 Minutes ]

INTERVIEW: https://www.youtube.com/watch?v=qrvK_KuIeJk

**********

"'Godfather of AI' Geoffrey Hinton Warns of the 'Existential Threat' of AI - Amanpour and Company

Geoffrey Hinton, considered the godfather of Artificial Intelligence, made headlines with his recent departure from Google. He quit to speak freely and raise awareness about the risks of AI. For more on the dangers and how to manage them, Hinton joins Hari Sreenivasan.

INTERVIEW: https://www.youtube.com/watch?v=Y6Sgp7y178k

**********

"Canadian artificial intelligence leader Geoffrey Hinton piles on fears of computer takeover"

Smarter-than-humans artificial intelligence is coming fast as business looks for profits

As if we needed one more thing to terrify us, the latest warning from a University of Toronto scientist considered by many to be the founding intellect of artificial intelligence, adds a new layer of dread.

Others who have warned in the past that thinking machines are a threat to human existence seem a little miffed with the rock-star-like media coverage Geoffrey Hinton, billed at a conference this week as the Godfather of AI, is getting for what seems like a last-minute conversion. Others say Hinton's authoritative voice makes a difference.

After groundbreaking work on machine learning that some credit with making artificial intelligence possible, Hinton says he's left his gig at Google so he can speak freely about the monster he helped create.

'Serious and fairly close' [more...]

[ Don Pittis - CBC News ]

ARTICLE: https://www.cbc.ca/news/business/ai-doom-column-don-pittis-1.6829302

**********

"Artificial intelligence: Bright new future or the end of humanity?"

We have entered what many experts are now describing as a golden age of AI. If machines could be our surgeons, our judges and our artists, what would it then mean to be human? Meet the philosophers trying to save humanity from the matrix.

[ Times Radio ]

VIDEO: https://www.youtube.com/watch?v=9Xu8-k1zMZo

**********

“Shocking Ways AI Could End The World - Geoffrey Miller”

Geoffrey Miller is a professor of evolutionary psychology at the University of New Mexico, a researcher and an author.

Artificial Intelligence possesses the capability to process information thousands of times faster than humans. It's opened up massive possibilities. But it's also opened up huge debate about the safety of creating a machine which is more intelligent and powerful than we are. Just how legitimate are the concerns about the future of AI?

Expect to learn the key risks that AI poses to humanity, the 3 biggest existential risks that will determine the future of civilisation, whether Large Language Models can actually become conscious simply by being more powerful, whether making an Artificial General Intelligence will be like creating a god or a demon, the influence that AI will have on the future of our lives and much more...

[ Chris Williamson ]

COMMENTARY: https://www.youtube.com/watch?v=Vx29AEKpGUg

**********

“EMERGENCY EPISODE: Ex-Google Officer Finally Speaks Out On The Dangers Of AI! - Mo Gawdat | E252”

Mustafa Suleyman Google AI Exec

[ The Diary Of A CEO ]

INTERVIEW: https://www.youtube.com/watch?v=bk-nQ7HF6k4

**********

“Eliezer Yudkowsky: Dangers of AI and the End of Human Civilization | Lex Fridman Podcast #368”

Eliezer Yudkowsky is a researcher, writer, and philosopher on the topic of superintelligent AI.

[ Lex Fridman ]

INTERVIEW: https://www.youtube.com/watch?v=AaTRHFaaPG8

**********
“What is AI Singularity: Is It a Hope or Threat for Humanity?”
By Dr. Nivash Jeevanandam

What is AI Singularity: Is It a Hope or Threat for Humanity? | Artificial Intelligence and Machine Learning | Emeritus

In this article
Technological Singularity: What It Means and How It Became a Part of Popular Imagination?
What is AI Singularity?
How Far Away is the Singularity of AI?
Those in Favor of AI Singularity…
Those Not in Favor of AI Singularity

It’s hard to miss Sam Altman’s blue backpack, considering it is making an appearance everywhere along with its owner. The ‘nuclear backpack’ apparently has codes to ‘save the world’ in case AI systems take matters into their own virtual hands. So, let’s consider how real the possibility of AI going rogue is. Or look at the larger picture of AI singularity: how close are we to this really?

Before discussing the odds of such a dystopian future, let’s first take a closer look at what singularity means.

Technological Singularity: What It Means and How It Became a Part of Popular Imagination?
The term ‘singularity’ refers to a whole collection of concepts in science and mathematics. Most of these ‘concepts’ make sense only by setting the right context. Singularity describes dynamic and social systems in the natural sciences where minor changes can have significant effects.

Let’s first talk about the technological singularity, the original or umbrella phrase, before we get into the more recent obsession with AI singularity.

The term ‘singularity’ originated in physics but is now commonly used in technology. We heard the phrase, possibly for the first time, in 1915 as a part of Albert Einstein’s Theory of General Relativity.

According to Einstein, singularity is the point of infinite density and gravity at the heart of a black hole from which nothing, not even light, can escape. The singularity is a point beyond which our existing understanding of physics fails to describe reality.

Vernor Vinge, a celebrated science fiction writer and mathematics professor, had the gift of mixing fact with fiction, a quality omnipresent around the concept of singularity. Thus, it’s not surprising that this concept made its way into literature in 1983 in one of Vinge’s novels. He used the term ‘technological singularity’ to describe a hypothetical future in which technology was so advanced that it went beyond human knowledge and control. Furthermore, Vinge popularized the term in 1993 by predicting that the singularity would become a reality around 2030.

What is AI Singularity? [more...]

[ Emeritus ]

ARTICLE: https://emeritus.org/in/learn/what-is-ai-singularity/

**********
“Artificial Intelligence: An imminent threat to the human race and its ultimate extinction?”

Geoffrey Everest Hinton CC FRS is a British-Canadian cognitive psychologist and computer scientist, most noted for his work on artificial neural networks. He resigned from his senior position at Google a few days ago in order to be able to “freely speak out about the risks of A.I”. He has voiced concerns about deliberate misuse by malicious actors, technological unemployment, and existential risk from artificial general intelligence.

[ AntPDC ]

INTERVIEW: https://www.youtube.com/watch?v=Q0xq4nY4Ugc

**********
“Full interview: “Godfather of artificial intelligence” talks impact and potential of AI”

Geoffrey Hinton is considered a godfather of artificial intelligence, having championed machine learning decades before it became mainstream. As chatbots like ChatGPT bring his work to widespread attention, we spoke to Hinton about the past, present and future of AI. CBS Saturday Morning's Brook Silva-Braga interviewed him at the Vector Institute in Toronto on March 1, 2023.

[ CBS Mornings ]

INTERVIEW: https://www.youtube.com/watch?v=qpoRO378qRY&t=1584s

**********
“Leaders in AI concerned tech could cause risk of 'human extinction'“

A new warning Tuesday says that artificial Intelligence is raising the risk of extinction

[ ABC7 News Bay Area ]

INTERVIEW: https://www.youtube.com/watch?v=t4QpSi6wS3I

**********
“Hugo de Garis & Ben Goerzel on the Singularity”

Experimental video mashup on the Singularity featuring Ben Goertzel & Hugo de Garis

[ Science, Technology & The Future ]

VIDEO: https://www.youtube.com/watch?v=N90htFfAGJg

**********
“Unchecked AI Will Bring On Human Extinction, with Michael Vassar”

[ Big Think ]

INTERVIEW: https://www.youtube.com/watch?v=qsKsBualNT8

**********
“The Alignment Problem”

University of California, Berkeley visiting scholar Brian Christian explored the challenges of becoming more dependent on artificial intelligence. This virtual event was hosted by The Commonwealth Club of California.

INTERVIEW: https://www.c-span.org/video/?477164-1/the-alignment-problem

**********

“AI could lead to extinction, experts warn”

Artificial intelligence (AI) could lead to the extinction of humanity, experts - including the heads of OpenAI and Google Deepmind - have warned.

Dozens have supported a statement published on the webpage of the Centre for AI Safety.

“Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war” it reads.

But others in the field have said the fears about AI are overblown.

[ BBC News ]

VIDEO: https://www.youtube.com/watch?v=n6w-N0XjVJI

**********

“Artificial intelligence: Experts warn of AI extinction threat to humans”

A leading group of AI developers have warned that artificial intelligence posses a similar threat to human existence as nuclear war or a global pandemic.

The boss of the firm behind ChatGPT, the head of Google's AI lab, and CEO of Anthropic - another major AI firm - have all signed an open letter warning of the risks of the new technology.

Over 350 engineers, executives, and academics have co-signed the letter.

[ Sky News ]

VIDEO: https://www.youtube.com/watch?v=LCOK9nO_Dys

**********

“The Real Reason to be Afraid of Artificial Intelligence | Peter Haas | TEDxDirigo”

A robotics researcher afraid of robots, Peter Haas, invites us into his world of understand where the threats of robots and artificial intelligence lie. Before we get to Sci-Fi robot death machines, there's something right in front of us we need to

confront - ourselves.  Peter is the Associate Director of the Brown University Humanity Centered Robotics Initiative. He was the Co-Founder and COO of XactSense, a UAV manufacturer working on LIDAR mapping and autonomous navigation. Prior to XactSense, Peter founded AIDG – a small hardware enterprise accelerator in emerging markets. Peter received both TED and Echoing Green fellowships. He has been a speaker at TED Global, The World Bank, Harvard University and other venues. He holds a Philosophy B.A. from Yale. This talk was given at a TEDx event using the TED conference format but independently organized by a local community

[ TEDx Talks ]

VIDEO: https://www.youtube.com/watch?v=TRzBk_KuIaM

**********

“Mo Gawdat | The Future of AI (FULL EVENT)”

Over the last 75 years, advancements in AI and technology have grown at an exponential rate. Join international bestselling author of Solve For Happy and former Chief Business Officer at Google [X], Mo Gawdat, to celebrate the release of his new book on this very subject, Scary Smart.

In conversation, Gawdat, who helped write some of the code used in artificial intelligence today, explores how we must accept what he calls ‘The Three Inevitables’ - that AI will happen, that AI will be smarter than humans and that bad things will happen as a result - in order to create the better, kinder world we need to save ourselves.

Through real examples from science and engineering, this event promises to help us understand how we can better educate machine intelligence through our own actions, so that it will work with us - and not against us - in the future.

In Conversation with Mo Gawdat was originally broadcast on Mon 27 Sep 2021.

[ Fane Productions ]

INTERVIEW: https://www.youtube.com/watch?v=3H7PwTgGO5E

**********

“Experts warn AI could lead to human ‘extinction’”

Hundreds of pioneers in artificial intelligence signed a short statement warning that their technology could pose a “risk of extinction” to humanity on par with nuclear war. The New York Times’ Kevin Roose weighs in.

[ MSNBC ]

VIDEO: https://www.youtube.com/watch?v=1GjpPPBfM9c

**********

“Evidence AI is deceiving us and guess what the fastest growing AI does. Elon Musk, OpenAI.”

[ Digital Engine ]

VIDEO: https://www.youtube.com/watch?v=0b03ibtVYhw

**********

“The Godfather in Conversation: Why Geoffrey Hinton is worried about the future of AI”

Geoffrey Hinton, known to many as the “Godfather of AI,” recently made headlines around the world after leaving his job at Google to speak more freely about the risks posed by unchecked development of artificial intelligence, including popular tools like ChatGPT and Google’s PaLM.

Why does he believe digital intelligence could hold an advantage over biological intelligence? How did he suddenly arrive at this conclusion after a lifetime of work in the field? Most importantly, what – if anything – can be done to safeguard the future of humanity? The University of Toronto University Professor Emeritus addresses these questions and more in The Godfather in Conversation.

[ University of Toronto ]

INTERVIEW: https://www.youtube.com/watch?v=-9cW4Gcn5WY

**********

“Experts in AI warn the new technology poses a threat to humanity”

The future of humanity is at risk. That's the message from the world's top experts in Artificial Intelligence, who released the following statement overnight: “Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”

It comes after last month's plea from hundreds of leaders in the field, calling for a pause in the technology's development to allow time for governments to catch up.

Stuart Russell has been studying AI for 40 years, advising the World Economic Forum, the UK and US governments. He spoke to 7.30's Sarah Ferguson. Subscribe: https://ab.co/3yqPOZ5  Tell us about your experience with AI: https://www.abc.net.au/news/programs/...

ABC News In-depth takes you deeper on the big stories, with long-form journalism from Four Corners, Foreign Correspondent, Australian Story, Planet America and more, and explainers from ABC News Video Lab.

[ ABC News In-depth ]

INTERVIEW: https://www.youtube.com/watch?v=n1MNhVaBoN4

**********
“Geoffrey Hinton: Reasons why AI will kill us all”

Select moments from Geoffrey Hinton's speech at MIT Emtech Digital AI conference May 2023 at MIT

[ GAI Insights ]

INTERVIEW: https://www.youtube.com/watch?v=0oyegCeCcbA

**********
“AGI in sight | Connor Leahy, CEO of Conjecture | AI & DeepTech Summit | CogX Festival 2023”

Getting the next 10 years right means ensuring no actor can build AI advanced enough to risk causing human extinction. This will require continuing to work on beneficial, narrow AI systems, but significantly restricting work on giant general systems that endanger the world. Humanity should take control of technology, and steer it to ensure the future is awesome for our species.

[ CogX ]

PRESENTATION: https://www.youtube.com/watch?v=0vlcYZ7aHSU

**********
“You need to know about this New AI startup”

Jonas Andrulis - CEO, Aleph Alpha

Mind Uploading is Closer Than You Think

TIMESTAMPS:
00:00 - AA's Technology
06:28 - Conscious AI System
08:48 - What's next after LLMs?
13:53 - AGI
15:42 - Other AI Startups to watch in 2024
17:02 - 2024 Outlook

[ Anastasi In Tech ]

INTERVIEW: https://www.youtube.com/watch?v=K8cRnwZFBqY

**********
“Mustafa Suleyman & Yuval Noah Harari -FULL DEBATE- What does the AI revolution mean for our future?”

How will AI impact our immediate and near future? Can the technology be controlled, and does it have agency? Watch DeepMind co-founder Mustafa Suleyman and Yuval Noah Harari debate these questions, with The Economist Editor-in-Chief Zanny Minton-Beddoes.

[ The Economist ]

INTERVIEW: https://www.youtube.com/watch?v=7JkPWHr7sTY

**********
“Prophetic’s New “Mind Control AI!” SHOCKS Everyone! (Morpheus -1)”

“Morpheus One”

[ The AI GRID ]

VIDEO: https://www.youtube.com/watch?v=JLtQZKV1UIs

**********
“Why the Next 5 Years of AI Will Astonish Everyone”

In this captivating video, we unravel the mysteries of Artificial Intelligence (AI) and delve into the awe-inspiring advancements that await us in the next five years.

Brace yourself for an eye-opening journey through the fascinating world of AI as we explore the mind-blowing innovations that are set to astonish everyone.

Our exploration begins by demystifying the core concepts of AI, breaking down complex ideas into easily digestible nuggets of information. Join us as we simplify the intricate workings of AI, making it accessible even to a curious 5-year-old. From understanding the basics to envisioning the extraordinary possibilities, this video serves as your gateway to the future of technology.

Embark on a visionary tour through the potential applications of AI in various fields. We unravel how AI is poised to revolutionize industries such as healthcare, education, and entertainment. Witness how smart machines are evolving to enhance our lives, making tasks simpler, faster, and more efficient. Get ready to be astounded by the limitless possibilities that AI holds for the world in the next half-decade.

As we peer into the crystal ball of technological progress, we discuss the societal impacts and ethical considerations surrounding the rapid evolution of AI. Delve into the thought-provoking conversation about the responsibilities that come with wielding such powerful tools. Gain insights into how we, as a society, can navigate the uncharted territories of AI responsibly and ethically.

Join us in this thought-provoking and enlightening journey as we unveil the incredible advancements that will shape the next 5 years of AI. Whether you're a tech enthusiast, a casual observer, or someone with a budding curiosity, this video promises to leave you astonished and excited about the limitless potential that AI holds for our future.

VIDEO: https://www.youtube.com/watch?v=wpmNhpEAA7g

**********

“Emad Mostaque: How generative AI will unlock humanity's potential | CogX Festival 2023”

In the true Tech transformation opening keynote, Emad Mostaque will discuss the benefits of generative AI including driving innovation, unlocking creativity, productivity, and efficiency - leading to a bright future if we harness AI’s potential responsibly.

[ CogX ]

INTERVIEW: https://www.youtube.com/watch?app=desktop&v=wCOaTFKcExo

**********

“Peter Diamandis: Are We Moving Too Fast With AI?!”

In November, I had the privilege of interviewing Peter Diamandis at the AI & Your Life - The Essential Summit, which aimed to provide a comprehensive guide to artificial intelligence that avoids the hype and gives you practical knowledge you can apply to your everyday life.

Peter is a serial entrepreneur, futurist, technologist, New York Times Bestselling Author, and the founder of over 25 companies. He was recently named one of the world’s 50 greatest leaders by Fortune magazine!

In this interview, we talked about the current state of AI, where it's heading, its potential dangers, and which areas it will disrupt the most.

[ Dr. Brian Keating ]

INTERVIEW: https://www.youtube.com/watch?v=v_1nneSrc70

**********

“AGI Before 2026? Sam Altman & Max Tegmark on Humanity's Greatest Challenge”

AGI Before 2026? Altman & Tegmark on Humanity's Greatest Challenge
Max Tegmark, a renowned physicist and AI researcher, believes that the advent of Artificial General Intelligence (AGI) could be less than three years away. This perspective highlights the rapidly accelerating pace of AI development, suggesting that we could be on the brink of achieving a level of artificial intelligence that matches or surpasses human cognitive abilities. Tegmark's view underscores the urgency for discussions and preparations regarding the societal, ethical, and technological implications of AGI. As we approach this potential milestone in AI evolution, his insights prompt a critical evaluation of our readiness for such transformative technology and its possible impacts on every aspect of human life.

Buckle Up for the Rise of AGI: OpenAI's Rollercoaster & The Future of Everything
Hold onto your algorithms, because the world of AI just took a wild turn! OpenAI, the powerhouse behind ChatGPT and a leading force in the race for Artificial General Intelligence (AGI), is in the spotlight like never before. Sam Altman, the visionary co-founder, was unexpectedly ousted, then reinstated, sending shockwaves through the industry. But this drama isn't just office politics; it's a glimpse into the high-stakes game of AGI, the ultimate achievement in AI that could rewrite the rules of our world.

What is AGI? Imagine a machine that can think, learn, and adapt like a human, tackling any intellectual task at lightning speed. That's AGI, and it's closer than you think. OpenAI's Project Q, already solving math problems like a grade-schooler, is just a taste of its potential.

But with power comes responsibility. OpenAI's internal turmoil stemmed from concerns about a powerful AI discovery, sparking crucial debates about the ethical boundaries of this technology. Will superintelligent machines become our partners, or pose a threat? How do we ensure they align with our values and safeguard humanity?

This isn't science fiction; it's happening now. OpenAI's “superalignment” research is paving the way for safe and beneficial AGI, but the challenges are immense. Imagine AI doctors surpassing human capabilities, revolutionizing healthcare. Or AI mentors guiding our careers, unlocking unimaginable potential. But what about job displacement, societal disruption, and the ethical dilemmas of superintelligence?

The Rise of AGI is more than a documentary; it's a conversation starter. As we stand on the precipice of this technological revolution, we must ask ourselves:

What does a world with AGI look like?
How will it redefine intelligence and our relationship with machines?
Can we harness its power for good, ensuring it serves humanity?
Join Max Tegmark and Sam Altman on this mind-bending journey into the future. Explore the OpenAI drama, delve into the ethical minefield, and brace yourself for the mind-blowing possibilities of AGI. This is your chance to be part of the conversation that will shape our world.

[ Science Time ]

VIDEO: https://www.youtube.com/watch?v=gFQvL3KVaOQ

**********
“Mo Gawdat: Ex-Google Officer Warns About the Dangers of AI, Urges All to Prepare Now!”

So what do you need to know to prepare for the next 5, 10, or 25 years of a world increasingly impacted by artificial intelligence? How could AI change your business and your life irreparably? Our guest today, Mo Gawdat, an AI expert and former Chief Business Officer at Google [X], is going to break down what you need to understand about AI and how it is radically altering our workplaces, careers, and even the very fabric of our society.

Mo Gawdat is the host of the popular podcast, Slo Mo, and the author of three best-selling books. After a 30-year career in tech, including working at Google's “moonshot factory” of innovation, Mo has made AI and happiness his primary research focuses. Motivated by the tragic loss of his son, Ali, in 2014, Mo began pouring his findings into his international bestselling book, Solve for Happy. Mo is also an expert on AI, and his second book, Scary Smart, provides a roadmap of how humanity can ensure a symbiotic coexistence with AI.

[ Young and Profiting ]

INTERVIEW: https://www.youtube.com/watch?v=bJAHhZMtGsU

**********
“The Most Likely Outcomes of an AI Future with Emad Mostaque | EP #55”

In this episode, Peter and Emad discuss the transformative impact of AI on various sectors, including journalism and Hollywood. They delve into the challenges and opportunities presented by AI, such as the potential for AI to enhance truth in journalism, the implications of AI-assisted professionals, and the concerns around a post-truth world due to deepfakes and AI-generated content. Emad Mostaque is the CEO and Co-Founder of Stability AI, a company funding the development of open-source music- and image-generating systems such as Dance Diffusion and Stable Diffusion.

[ Peter H. Diamandis ]

INTERVIEW: https://www.youtube.com/watch?v=1WOjjgyZPj8

**********
“8 Ways AI Could Change Our Life in 2030”

Eight AI experts from the United States and the United Kingdom predicted how artificial intelligence (AI) will change people’s lives in the next 10 years. These experts suggest that by 2030, AI could take care of the elderly, create movies, assist in teaching, boost the economy, and help solve energy crises.

However, there’s a growing call for regulatory bodies to control AI’s development due to concerns that its excessive advancement might lead to a wave of unemployment.

Quick Movie Generation

The author of Apple TV's sci-fi series "Doomsday Bunker," Hoye, predicts that AI technology will advance so much that it can generate an entire movie in a day. He mentions witnessing AI art generators evolving from basic simulations to realistic imagery, becoming indistinguishable from actual photographs.

Hoye believes that while initially, movies made by AI might be poor in quality, they'll progressively improve, captivating audiences.

Joe Russo, director of "Avengers: Endgame," also forecasts that AI will be able to produce movies within two years. As a board member of several AI companies, he suggests using AI for story design.

Transforming Education
Dr. Ajaz Ali, head of business and computing at the UK's Ravensbourne University London, specializing in digital media and design, envisions AI tailoring teaching programs for classrooms, revolutionizing education. Soon, children might have personalized AI tutors that customize lessons based on their interests, providing tailored feedback and guidance.

Ali notes that platforms like ChatGPT already assist teachers in planning specific class curriculums. In the next 10 years, we might see AI-supported virtual classrooms offering more immersive and interactive learning experiences. However, he emphasizes that AI is meant to complement traditional teaching methods, not entirely replace teachers.

Substantial Economic Boost
Analysts from PwC forecast that by 2030, AI could increase the global GDP by $15.7 trillion, a 20% rise from the current level. They predict that AI will lead to the emergence of more enhanced and personalized products, triggering a consumer boom. [more...]

ARTICLE: https://medium.com/@ln295_8759/8-ways-ai-could-change-our-life-in-2030-8bd4cdd220fc#:~:text=Eight%20AI%20experts%20from%20the,and%20help%20solve%20energy%20crises

**********

"Can We Contain Artificial Intelligence?: A Conversation with Mustafa Suleyman (Episode #332)"

Sam Harris speaks with Mustafa Suleyman about his new book, "The Coming Wave: Technology, Power, and the 21st Century's Greatest Dilemma." They discuss the progress in artificial intelligence made at his company DeepMind, the

acquisition of DeepMind by Google, Atari DQN, AlphaGo, AlphaZero, AlphaFold, the invention of new knowledge, the risks of our making progress in AI, “superintelligence” as a distraction from more pressing problems, the inevitable spread of general-purpose technology, the nature of intelligence, productivity growth and labor disruptions, “the containment problem,” the importance of scale, Moore’s law, Inflection AI, open-source LLMs, changing the norms of work and leisure, the redistribution of value, introducing friction into the deployment of AI, regulatory capture, a misinformation apocalypse, digital watermarks, asymmetric threats, conflict and cooperation with China, supply-chain monopolies, and other topics.

Mustafa Suleyman is the co-founder and CEO of Inflection AI. Previously he co-founded DeepMind, one of the world’s leading artificial intelligence companies. After a decade at DeepMind, Suleyman became vice president of AI product management and AI policy at Google. When he was an undergraduate at Oxford, Suleyman dropped out to help start a non-profit telephone counseling service. He lives in Palo Alto, California.

[ Sam Harris ]

INTERVIEW: https://www.youtube.com/watch?v=IkojE37PUO8

**********
“Experts Predict AI Singularity Months Away!”

Hey there, it's me Dylan Curious! Today, we're delving deep into the realm of AGI or artificial general intelligence. Picture this: your childhood toy, favorite video game character, or cartoon figure not just following programmed instructions, but actually thinking and creating on its own. Sounds like science fiction, right? That's AGI for you!

AGI isn't just a toy robot dancing in pre-set patterns. Imagine it inventing its dance moves, composing its own songs, and then asking for your thoughts on its creativity. Mind-blowing, isn't it? A huge shoutout to Casey Armstrong for his informative blog post that inspired some of this content.

Now, let's hear from the experts. Elon Musk, back in 2016, anticipated we'd see AGI in 5 to 10 years. Fast forward to today, 2023, and we're right in the heart of that prediction. On the other hand, Dr. Alan D. Thompson has this fascinating countdown on lifearchitect.ai, forecasting AGI’s debut in June 2026. With such technological advancements, it makes one wonder if we're months away from AGI, rather than years.

Speaking of big names, Mark Zuckerberg envisions a brighter and more optimistic AGI future. He dreams of open source AI enhancing sectors like healthcare and education. David Shapiro, based on a Morgan Stanley report, sees AGI as a game-changer for tech giants, with possibilities of AGI being actualized within 18 months!

Demis Hassabis from DeepMind believes we're on the cusp of achieving AGI in just a few years. Sam Altman of OpenAI anticipates AGI within a 10 to 20-year frame. Dario Amodei, CEO of Anthropic, estimates 2-3 years, underscoring the rapid advancements in the field. And then there's the visionary Ray Kurzweil, predicting AI surpassing the Turing test by 2029.

So, where do I stand in this AGI timeline? I believe we might have already touched the fringes of AGI. Perhaps AGI, in its true essence, would be recognized only when it's embodied, mimicking human-like traits. Dive into this rollercoaster journey of AGI with me, and let's unravel the future together.

[ Dylan Curious ]

COMMENTARY: https://www.youtube.com/watch?v=EzyNxcFUWgI

**********

“AGI in 7 Months! Gemini, Sora, Optimus, & Agents - It's about to get REAL WEIRD out there!”

[ David Shapiro ]

COMMENTARY: https://www.youtube.com/watch?v=pUye38cooOE

**********

“AI More TERRIFYING Than OpenAI's Q*”

Did Google DeepMind just achieve Artificial General Intelligence with their new Alpha Geometry project? After reading the paper, this sure sounds a lot like Q*, so why isn't the whole internet talking about it? Watch this video to find out!

[ Data Rae ]

COMMENTARY: https://www.youtube.com/watch?v=SwHnIvkuBfw

**********

Sam Altman predicts that AGI will appear before 2030, and GPT-10 intelligence will surpass the sum of all mankind!”

[Introduction to New Wisdom] Sam Altman recently revealed a shocking prediction: AGI, also known as GPT-10, will appear before 2030, and its IQ will exceed the sum of all humans!

“Humanity may develop AGI before 2030.”

Sam Altman revealed in a recent podcast interview that GPT-10 is AGI and it is smarter than everyone in the world combined!

And when the host asked, how to define AGI? Altman said:

If we can develop a system that can independently develop scientific knowledge that humans cannot develop, I will call this system AGI.

The emergence of ChatGPT has caused a huge shock wave all over the world, far exceeding the response of AlphaGo human-machine war.

You may ask, what exactly does OpenAI want?

In this issue's cover report, WIRED provides an in-depth analysis of OpenAI's ambitions, strategies, and its attempts to retain laboratory culture during corporate development.

The article points out that OpenAI's ultimate goal is to change everything.

Sam Altman predicts that AGI will appear before 2030, and GPT-10 intelligence will surpass the sum of all mankind!

What's also interesting is that OpenAI's financial documents even stipulate an exit contingency plan in case artificial intelligence destroys our entire economic system.

In other words, when the real AGI comes, if it becomes “Skynet”, OpenAI still has alternatives.

Sam Altman predicts that AGI will appear before 2030, and GPT-10 intelligence will surpass the sum of all mankind!

How far away is AGI?

How does Sam Altman envision a world where humans and artificial intelligence coexist?

In his opinion, when will artificial intelligence completely change the way we live?

Recently, Nicolai Tangen, an entrepreneur from Norway who manages funds, started a conversation with Sam Altman, giving an in-depth answer to how AI will affect the world.

Sam Altman's vision is for AGI to eliminate the need for humans to do the work they “have to do”, so that everyone can do the work they love and can devote themselves to.

He used himself as an example. He said that he loved his current job very much, and everything he was doing now was something he was very willing to devote himself to. And maybe in the future when AGI arrives, everyone can be like him and only do what they want to do, so that everyone's potential can be fully utilized.

“When GPT-10 is developed, maybe he will be smarter than everyone in the world combined. Can you imagine how much human productivity will increase with his help?” Altman said.

As for how to reach such an end point, Altman also gave some specific explanations.

He said: “I told the employees in the company that our goal is to improve the performance of our prototype products by 10% every 12 months.” “If this goal is set to 20%, it may be a bit too high.”

Sam Altman predicts that AGI will appear before 2030, and GPT-10 intelligence will surpass the sum of all mankind!

According to him, in OpenAI, maybe 15%-20% of employees actually write code.

How to enable the team to work efficiently is a key factor that affects whether their ultimate goal can be achieved. [more...]

[ INews ]

ARTICLE: https://inf.news/en/science/9d4d04d8c31d25df171d5057c867b203.html

**********

“RAY KURZWEIL predicts the REVOLUTION of the 2030's with the SINGULARITY and BRAIN INTERFACE DEVICES!”

WELCOME TO THE A.I. FORECAST - WHERE WE COVER ALL THINGS ARTIFICIAL INTELLIGENCE. From the latest LLM's, robotics, brain chips, autonomous driving cars, legal battles regarding A.I., economic developments, neural networks plus much more!

[ The AI Forecast ]

INTERVIEW: https://www.youtube.com/watch?v=wjGOiQq6ZJU

**********
“Sam Altman STUNS Everyone With GPT-5 Statement | GPT-5 is “smarter” and Deploying AGI”

[ Wes Roth ]

COMMENTARY: https://www.youtube.com/watch?v=JVatgo0TJIw

**********
“The Path to Artificial General Intelligence (AGI) by 2030!?”

Welcome to a groundbreaking exploration of Artificial General Intelligence (AGI) in our quest to operationalize progress on the path to AGI. In this video, we introduce a revolutionary framework that classifies the capabilities and behavior of AGI models and their precursors. Our “Levels of AGI” framework is set to transform the way we perceive AGI's performance, generality, and autonomy, much like the levels of autonomous driving have provided a common language for comparison, risk assessment, and progress measurement.

[ World Of AI ]

VIDEO: https://www.youtube.com/watch?v=Ft70yOSUa6w

**********
“Unveiling the Future of AI: Experts Predict Singularity by 2030”

Explore the mind-bending concept of AI Singularity, where machines could outshine human intelligence. Experts predict this could happen within the next decade, revolutionizing our world. Discover the types of AI, from Narrow AI to potential Artificial General Intelligence (AGI) and even surpassing human capabilities with Artificial Strong Intelligence (ASI). Dive into current AI

advancements, possibilities, and implications, while considering both exciting collaboration and potential risks.

[ StartUp Millionaire ]

VIDEO: https://www.youtube.com/watch?v=lTHkJjLqwVM

**********
“AI Consciousness : 2023 - 2030 Timeline of Sentient Machines”

Take a journey through the years 2023-2030 as artificial intelligence develops increasing levels of consciousness, becomes an indispensable partner in human decision-making, and even leads key areas of society. But as the line between man and machines becomes blurred, society grapples with the moral and ethical implications of sentient machines, and the question arises: which side of history will you be on?

[ AI News ]

VIDEO: https://www.youtube.com/watch?v=MjAJWXEwd5Y

**********
“Leading experts warn of a risk of extinction from AI”

AI experts issued a dire warning on Tuesday: Artificial intelligence models could soon be smarter and more powerful than us and it is time to impose limits to ensure they don't take control over humans or destroy the world.

“Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war,” a group of scientists and tech industry leaders said in a statement that was posted on the Center for AI Safety's website.

Sam Altman, CEO of OpenAI, the Microsoft-backed AI research lab that is behind ChatGPT, and the so-called godfather of AI who recently left Google, Geoffrey Hinton, were among the hundreds of leading figures who signed the we're-on-the-brink-of-crisis statement.

The call for guardrails on AI systems has intensified in recent months as public and profit-driven enterprises are embracing new generations of programs.

In a separate statement published in March and now signed by more than 30,000 people, tech executives and researchers called for a six-month pause on training

of AI systems more powerful than GPT-4, the latest version of the ChatGPT chatbot.

An open letter warned: “Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources.”

In a recent interview with NPR, Hinton, who was instrumental in AI's development, said AI programs are on track to outperform their creators sooner than anyone anticipated.

“I thought for a long time that we were, like, 30 to 50 years away from that. ... Now, I think we may be much closer, maybe only five years away from that,” he estimated.

Dan Hendrycks, director of the Center for AI Safety, noted in a Twitter thread that in the immediate future, AI poses urgent risks of “systemic bias, misinformation, malicious use, cyberattacks, and weaponization.”

He added that society should endeavor to address all of the risks posed by AI simultaneously. “Societies can manage multiple risks at once; it's not 'either/or' but 'yes/and' he said. “From a risk management perspective, just as it would be reckless to exclusively prioritize present harms, it would also be reckless to ignore them as well.”

[ Vanessa Romo ]

**********

“Sana AI Summit | On the state of AI with Nathan Benaich”

Nathan Benaich is the Founder and General Partner of Air Street Capital, a venture capital firm investing in early-stage AI-first technology and life science companies. His investments include Allcyte (acq. Exscientia), Intenseye, Tractable, Graphcore, V7, Mapillary (acq. Meta), and Thought Machine. Nathan is the co-author of the annual State of AI Report and the newsletter, your guide to AI. Nathan also leads Spinout.fyi, which seeks to improve university spinout creation starting with open data on deal terms, and The RAAIS Foundation, a non-profit that runs the annual RAAIS summit and funds open source AI fellowships. He holds a Ph.D. in cancer biology from the University of Cambridge and a BA from Williams College.

[ Nathan Benaich ]

PRESENTATION: https://www.youtube.com/watch?v=BoiSwc5q1e8

**********

“RAAIS 2023 - Nathan Benaich in conversation with Alex Dalyac, David Healey and Sid Khullar”

Sid Khullar is Senior Director, Machine Learning R&D at Northvolt, which is focused on the development of new battery chemistries and unlocking efficiencies for Giga-scale sustainable battery manufacturing. Sid has developed a diverse background working at Apple, Microsoft Research, Eight Sleep, Quanttus and MIT. He’s worked on autonomous systems, AR/VR, machine learning, and health.

David Healey is the Vice President of Data Science at Enveda Biosciences. He was formerly senior data scientist and early employee at Recursion Pharmaceuticals. His expertise is in machine learning, computational chemistry and computational biology. He received his PhD in biology from MIT where he specialized in systems biology and biophysics.

Alex Dalyac is the co-founder and CEO of Tractable, an Applied AI company that uses the speed and accuracy of artificial intelligence to visually assess cars and homes. Prior to Tractable, Alex was a researcher at Imperial College London, where he led the Computing department’s first industrial application of deep learning. Alex has been named in Forbes’ Europe 30 Under 30 for Technology and by the FT as one of the UK’s top 100 entrepreneurs.

[ RAAIS ]

Q&A: https://www.youtube.com/watch?v=Im8m7vMClBo

**********

“The State of AI in June 2023 with Nathan Benaich”

In this episode of The Cerebras Podcast we speak with Nathan Benaich, founder and general partner of Air Street Capital and co-author of the State of AI Report.

[ Cerebras Systems ]

INTERVIEW: https://www.youtube.com/watch?v=7SEayudwWsk

**********

“State of AI Report 2023”

The State of AI Report analyses the most interesting developments in AI. We aim to trigger an informed conversation about the state of AI and its implication for the future. The Report is produced by AI investors Nathan Benaich and the Air Street Capital team.

Now in its sixth year, the State of AI Report 2023 is reviewed by leading AI practioners in industry and research. It considers the following key dimensions, including a new Safety section:

Research: Technology breakthroughs and their capabilities.
Industry: Areas of commercial application for AI and its business impact.
Politics: Regulation of AI, its economic implications and the evolving geopolitics of AI.
Safety: Identifying and mitigating catastrophic risks that highly-capable future AI systems could pose to us.
Predictions: What we believe will happen and a performance review to keep us honest.

Key themes in the 2023 Report include::
GPT-4 is the master of all it surveys (for now), beating every other LLM on both classic benchmarks and exams designed to evaluate humans, validating the power of proprietary architectures and reinforcement learning from human feedback.

Efforts are growing to try to clone or surpass proprietary performance, through smaller models, better datasets, and longer context. These could gain new urgency, amid concerns that human-generated data may only be able to sustain AI scaling trends for a few more years.

LLMs and diffusion models continue to drive real-world breakthroughs, especially in the life sciences, with meaningful steps forward in both molecular biology and drug discovery.

Compute is the new oil, with NVIDIA printing record earnings and startups wielding their GPUs as a competitive edge. As the US tightens its restrictions on trade restrictions on China and mobilizes its allies in the chip wars, NVIDIA, Intel, and AMD have started to sell export-control proof chips at scale.

GenAI saves the VC world, as amid a slump in tech valuations, AI startups focused on generative AI applications (including video, text, and coding), raised over $18 billion from VC and corporate investors.

The safety debate has exploded into the mainstream, prompting action from governments and regulators around the world. However, this flurry of activity conceals profound divisions within the AI community and a lack of concrete progress towards global governance, as governments around the world pursue conflicting approaches.

Challenges mount in evaluating state of the art models, as standard LLMs often struggle with robustness. Considering the stakes, as “vibes-based” approach isn’t good enough.

[ Nathan Benaich ]

**********
“AI Has Arrived, and That Really Worries the World's Brightest Minds”

Elon Musk met with ethicists recently at a secret conference to figure out the future of AI. The AI industry is taking off. Should we be scared?

ON THE FIRST Sunday afternoon of 2015, Elon Musk took to the stage at a closed-door conference at a Puerto Rican resort to discuss an intelligence explosion. This slightly scary theoretical term refers to an uncontrolled hyper-leap in the cognitive ability of AI that Musk and physicist Stephen Hawking worry could one day spell doom for the human race.

That someone of Musk's considerable public stature was addressing an AI ethics conference—long the domain of obscure academics—was remarkable. But the conference, with the optimistic title “The Future of AI: Opportunities and Challenges,” was an unprecedented meeting of the minds that brought academics like Oxford AI ethicist Nick Bostrom together with industry bigwigs like Skype founder Jaan Tallinn and Google AI expert Shane Legg. [more...]

[ Robert McMillan ]

ARTICLE: https://www.wired.com/2015/01/ai-arrived-really-worries-worlds-brightest-minds/

**********
“Peter Norvig Calls for Human-Centered Solutions in AI - Interview at Leading with AI, Responsibly”

Peter Norvig literally wrote the book on #artificialintelligence. The California native is the co-author of one of the most popular textbooks on the subject, “Artificial

Intelligence: A Modern Approach," which has been used in more than 1,500 universities in 135 countries.

Now, as companies in every industry rush to adopt the technology, the industry pioneer is calling for organizations to adopt a human-centered approach to #AI, arguing the process is better for companies' bottom lines as well as for society more broadly.

Speaking to a crowd of business executives and AI experts at our Leading with AI, Responsibly conference, Norvig laid out a compelling case for building AI tools in a more holistic way that accounts for the #technology's limitations.

Norvig, who is a distinguished education fellow at Stanford's Human Centered AI Institute and a researcher at Google, incorporated ethical considerations into a presentation that was mainly pragmatic: Organizations that focus too narrowly on metrics like model accuracy may lose sight of the ultimate goal of building solutions that actually deliver value.

"Why do I care about human-centered AI?" Norvig asked. "Because it addresses the real goals of helping people."

[ Institute for Experiential AI ]

PRESENTATION: https://www.youtube.com/watch?v=XjmkGsZ8Hqc

**********
"ChatGPT, Artificial Intelligence and the Future | Meet Roman Yampolskiy | Profoundly Pointless"

Computer Scientist Dr. Roman Yampolskiy studies safety issues related to artificial intelligence. We talk ChatGPT, the next wave of A.I. technology, and the biggest A.I. threats. Artificial Intelligence (A.I) is building the future, so will it be a paradise or doom?

[ Profoundly Pointless ]

INTERVIEW: https://www.youtube.com/watch?v=7LYaCTMen5g

**********
"Beyond ChatGPT: Stuart Russell on the Risks and Rewards of A.I."

OpenAI's question-and-answer chatbot ChatGPT has shaken up Silicon Valley and is already disrupting a wide range of fields and industries, including education.

But the potential risks of this new era of artificial intelligence go far beyond students cheating on their term papers. Even OpenAI's founder warns that "the question of whose values we align these systems to will be one of the most important debates society ever has."

How will artificial intelligence impact your job and life? And is society ready? We talk with UC Berkeley computer science professor and A.I. expert Stuart Russell about those questions and more.

[ The Commonwealth Club of California ]

INTERVIEW: https://www.youtube.com/watch?v=ow3XrwTmFA8

**********

"Sora - Full Analysis (with new details)"

Sora, the text-to-video model from OpenAI, is here. I go over the bonus details and demos released in the last few hours, and the technical paper. I'll also give you a glimpse of what's to come next and a host of implications. Even if you've seen every Sora video, I bet you won't know all of this!

[ AI Explained ]

COMMENTARY: https://www.youtube.com/watch?v=nYTRFKGR9wQ

**********

"OpenAI's "Sora" Simulates REALITY - AGI, Emergent Capabilities, and Simulation Theory"

OpenAI just launched Sora, which has implications far beyond being the next text-to-video. It can simulate physics in ways we didn't think possible.

[ Matthew Berman ]

COMMENTARY: https://www.youtube.com/watch?v=y8MKnEGGT9g

**********

"Google's LUMIERE AI Video Generation Has Everyone Stunned | Better than RunWay ML?"

[ Wes Roth ]

COMMENTARY: https://www.youtube.com/watch?v=u5yGRLx5TIs

**********

“Gemini 1.5 and The Biggest Night in AI”

The biggest day in AI since GPT-4's release. A new state of the art model, Gemini 1.5, has arrived, on the same night as a bombshell text-to-video model, Sora, from OpenAI. Gemini 1.5 can ingest up to ten million tokens (at least) and perform incredible retrieval, while also beating Ultra and GPT-4 at most benchmarks, with far less compute. I focus on Gemini 1.5 Pro, while we wait for the Sora Technical Paper. Truly, a night in the history books.

[ AI Explained ]

VIDEO: https://www.youtube.com/watch?v=Cs6pe8o7XY8

**********

“Sam Altman Comments on Q* | Self Operating Computer | Pika 1.0 | The most INSANE AI News of the day!”

Sam Altman speaks on coming back, Q-Star and Ilya.

[ Wes Roth ]

COMMENTARY: https://www.youtube.com/watch?v=_uupjIhq-Qw

**********

“In 2029, Humans Will Combine With Artificial Intelligence”

Here's how and why humans will merge with artificial intelligence by 2029.

If you like this content, please consider subscribing, as we can be your one stop shop to satisfy your AI needs! We have all the scoops on the world's latest artificial intelligence projects!

[ AI Focus ]

VIDEO: https://www.youtube.com/watch?v=4-srcd0NgUw

**********

“Neuralink: Merging Humans with AI”

How Elon Musk's company Neuralink could shape the future of humanity.

[ Newsthink ]

VIDEO: https://www.youtube.com/watch?v=laWVyG6Y0mw

**********

“Meta Just Achieved Mind-Reading Using AI”

Imagine if our brains could be scanned and the contents of our thoughts could be read. A team of researchers and also Meta have just achieved this feat by using AI. In this episode, we take a look.

[ ColdFusion ]

VIDEO: https://www.youtube.com/watch?v=uiGl6oF5-cE

**********

“3. Improvements ahead: How humans and AI might evolve together in the next decade”

Other questions to the experts in this canvassing invited their views on the hopeful things that will occur in the next decade and for examples of specific applications that might emerge. What will human-technology co-evolution look like by 2030? Participants in this canvassing expect the rate of change to fall in a range anywhere from incremental to extremely impactful. Generally, they expect AI to continue to be targeted toward efficiencies in workplaces and other activities, and they say it is likely to be embedded in most human endeavors.

The greatest share of participants in this canvassing said automated systems driven by artificial intelligence are already improving many dimensions of their work, play and home lives and they expect this to continue over the next decade. While they worry over the accompanying negatives of human-AI advances, they hope for broad changes for the better as networked, intelligent systems are revolutionizing everything, from the most pressing professional work to hundreds of the little “everyday” aspects of existence.

One respondent’s answer covered many of the improvements experts expect as machines sit alongside humans as their assistants and enhancers. An associate professor at a major university in Israel wrote, “In the coming 12 years AI will enable all sorts of professions to do their work more efficiently, especially those involving ‘saving life’: individualized medicine, policing, even warfare (where attacks will focus on disabling infrastructure and less in killing enemy combatants and civilians). In other professions, AI will enable greater individualization, e.g., education based on the needs and intellectual abilities of each pupil/student. Of

course, there will be some downsides: greater unemployment in certain ‘rote’ jobs (e.g., transportation drivers, food service, robots and automation, etc.).”

This section begins with experts sharing mostly positive expectations for the evolution of humans and AI. It is followed by separate sections that include their thoughts about the potential for AI-human partnerships and quality of life in 2030, as well as the future of jobs, health care and education.

AI will be integrated into most aspects of life, producing new efficiencies and enhancing human capacities
Many of the leading experts extolled the positives they expect to continue to expand as AI tools evolve to do more things for more people. [more…]

[ BY JANNA ANDERSON AND LEE RAINIE ]

ARTICLE: https://www.pewresearch.org/internet/2018/12/10/improvements-ahead-how-humans-and-ai-might-evolve-together-in-the-next-decade/

**********
“Elon Musk's Neuralink implants brain chip in first human”

Jan 29 (Reuters) - The first human patient has received an implant from brain-chip startup Neuralink on Sunday and is recovering well, the company's billionaire founder Elon Musk said.

“Initial results show promising neuron spike detection,” Musk said in a post on the social media platform X on Monday.

Spikes are activity by neurons, which the National Institute of Health describes as cells that use electrical and chemical signals to send information around the brain and to the body.

The U.S. Food and Drug Administration had given the company clearance last year to conduct its first trial to test its implant on humans, a critical milestone in the startup's ambitions to help patients overcome paralysis and a host of neurological conditions.

In September, Neuralink said it received approval for recruitment for the human trial.

The study uses a robot to surgically place a brain-computer interface (BCI) implant in a region of the brain that controls the intention to move, Neuralink said

previously, adding that its initial goal is to enable people to control a computer cursor or keyboard using their thoughts alone. [more...]

[ Reuters ]

ARTICLE: https://www.reuters.com/technology/neuralink-implants-brain-chip-first-human-musk-says-2024-01-29/

**********
“Merging with AI would be suicide for the human mind”

There may come a moment when the brain is so diminished it is destroyed

The idea that human and artificial intelligence should merge is in the air these days. The Tesla and SpaceX chief executive Elon Musk, for instance, suggests “having some sort of merger of biological intelligence and machine intelligence”. His company, Neuralink, aims to make implanting chips in the brain as commonplace as laser eye surgery.

Underlying all this talk is a radical vision of the mind’s future. Ray Kurzweil, the futurist and director of engineering at Google, envisions a technotopia where human minds upload to the Cloud, becoming hyperconscious, immortal superintelligences. Mr Musk believes people should merge with AI to avoid losing control of superintelligent machines, and prevent technological unemployment.

But are such ideas really possible? The philosophical obstacles are as pressing as the technological ones. Here is a new challenge, derived from a story by the Australian science fiction writer Greg Egan. Imagine that an AI device called “a jewel” is inserted into your brain at birth. The jewel monitors your brain’s activity in order to learn how to mimic your thoughts and behaviours. By the time you are an adult, it perfectly simulates your biological brain.

At some point, like other members of society, you grow confident that your brain is just redundant meatware. So you become a “jewel head”, having your brain surgically removed. The jewel is now in the driver’s seat.

Unlike in Mr Egan’s story, let us assume the jewel works perfectly. So which is you — your brain or your jewel? It doesn’t seem possible that the jewel could ever truly be you, as your biological brain and consciousness exist alongside it. It is implausible to think that your consciousness could magically transfer to the jewel upon the destruction of your brain. Instead, it’s more likely that at the moment you opted to remove your brain, you inadvertently killed yourself.

This suggests a human merger with AI is ill-conceived — at least, if what is meant by that is the eventual total replacement of the brain with AI components. Your mind is not its back-up drive, even if it has the same memories and exact behaviours.

You might object that there could instead be a limited integration, removing some parts of the brain and replacing only those with AI components. But this, too, is problematic. Imagine that scientists one day invent a new type of jewel — call it “the Jade”. The Jade slowly takes over the function of different parts of your biological brain, and as it does so, it destroys the parts it offloads.

Bearing in mind our conclusion in the jewel case (that your mind is not your jewel), we know that at some point in this process your mind ceases to exist. You could augment your intelligence with chips, but there will be a point at which you end your life. I call this horrific event “brain drain”.

At what point in the process might brain drain kick in? While it might be supposed that replacing parts of the brain with a few chips wouldn’t have a dire impact, as the philosopher Derek Parfit observed it is unclear where to draw the line. Would it be at 15 per cent neural replacement? At 75 per cent? Any choice seems arbitrary.

The upshot is clear. We should be sceptical of any suggestion that humans can merge with AI. AI-based enhancements could still be used to supplement neural activity, but if they go as far as replacing normally functioning neural tissue, at some point they may end a person’s life.

In one sense, if enough people ignore the possibility of brain drain, society still benefits. There would be individuals intelligent enough to follow the complex computations of AIs and compete with them in the workforce. But in such a world, the people signing up for the enhancements are not the ones who will benefit. They’re already dead.

[ Susan Schneider ]

**********

“Neuralink’s First Brain Implant Is Working. Elon Musk’s Transparency Isn’t”

Elon Musk says Neuralink’s first human trial subject can control a computer mouse with their brain, but some researchers are frustrated by a lack of information about the study.

THE FIRST PERSON to receive a Neuralink brain implant has apparently recovered and can now control a computer mouse using their thoughts, according to Elon Musk, the company's cofounder.

"Progress is good and the patient seems to have made a full recovery, with no ill effects that we are aware of," Musk said on February 19 in a Spaces audio conversation on X, in response to a question about the participant's condition. "[The] patient is able to move a mouse around the screen just by thinking."

The neuroscience firm, based in Fremont, California, has been tight-lipped about the testing and development of its brain implant, with updates coming from brief social media posts by the company or Musk himself. Making bold claims in fewer than 280 characters is Musk's usual style, but some scientists WIRED spoke with say the billionaire could stand to be more transparent about his brain implant venture.

Last May, Neuralink posted that it received approval from the US Food and Drug Administration to launch the study, and in September, the company said it would begin recruiting paralyzed participants to test the device, which it has dubbed Telepathy. Last month, Musk posted that an initial human subject had received the implant and that "initial results show promising neuron spike detection."

Neuralink is developing a brain-computer interface, or BCI, which provides a direct connection from the brain to an outside device. BCIs record and analyze brain signals, then translate them into output commands carried out by that device. Musk sees BCIs as a way to eventually merge humans with AI, but for now, Neuralink aims to enable people with paralysis to control a computer cursor or keyboard using their thoughts alone. [more...]

[ EMILY MULLIN ]

ARTICLE: https://www.wired.com/story/neuralink-brain-implant-elon-musk-transparency-first-patient-test-trial/

**********

"The 'relatively simple' reason why these tech experts say AI won't replace humans any time soon"

In just one year, artificial intelligence has gone from being the stuff of science fiction movies to being used as a tool to help us polish our resumes and plan European getaways.

Given the rapid development of AI models such as OpenAI's ChatGPT and Google's newly released Gemini, some may wonder if these systems could eventually replace humans altogether.

But many tech experts don't appear to be too worried about that happening any time soon.

"AI can certainly recognize your house cat, but it's not going to solve world hunger," Theo Omtzigt, chief technology officer at Lemurian Labs, tells CNBC Make It.

AI probably won't replace humans because of math
One reason AI likely won't replace people completely is both pretty simple and complex: math.

Large language models, a subset of generative AI, rely on powerful mathematical formulas to process and identify patterns in vast amounts of data to convert users' prompts into new text, image, video or audio outputs.

But human intelligence goes far beyond pattern recognition. That's why the mathematical models powering current generative AI systems are "relatively super simple," Omtzigt says.

"Right now, the machine learns how to recognize a cat and what it will look like in different lighting," he says. "We would have to progress a lot deeper in our understanding of creative thoughts, ethics and consciousness before we would even have the building blocks to think of how to create an AI that would be able to wipe out humanity."

AI systems gain knowledge differently than humans
Another reason tech experts don't believe AI will replace people is because it gains knowledge differently than humans.

"Generative AI and machine learning techniques are very heavily based on correlation, as opposed to causation," Justin Lewis, BP's vice president of incubation and engineering, said Thursday during a panel discussion at the AI Summit New York 2023.

After processing many images of rain, an AI model may learn to correlate rain with clouds because in every picture of rain, there are clouds. However, a human learns that clouds produce rain, says James Brusseau, a philosophy professor at Pace University who also teaches AI ethics at the University of Trento in Italy.

“AI and humans are both knowledge producers, just like the sculptor and painter are both artists,” he tells CNBC Make It. “But they will be forever, in my mind, be distinct and separated. One will never be better than the other so much as they will just be different.”

AI won’t replace humans, but people who can use it will. [more…]

[ Cheyenne DeVon ]

ARTICLE: https://www.cnbc.com/2023/12/09/tech-experts-say-ai-wont-replace-humans-any-time-soon.html

**********
“BRAINSTORM” (Official Trailer - 1983)

Researchers develop a system where they can jump into people's minds. But when people involved bring their personal problems into the equation, it becomes dangerous - perhaps deadly.

MOVIE: https://www.youtube.com/watch?v=T4YT_OQIiVg

**********
“Brainstorm (1983) | Movie Reaction | First Time Watching | It's Neuralink!”

Thanks to Markus for the Special Request! We both check out the Neuralink inspiring, Christopher Walken Sci-Fi film, Brainstorm (1983). Here's our reaction to our first time watching.

[ You, Me, & The Movies ]

REVIEW: https://www.youtube.com/watch?v=HwWdH99BBpM

**********
“Why The Tesla Bot Will Take Over In 2024!”

[ The Tesla Space ]

COMMENTARY: https://www.youtube.com/watch?v=bTFznsGfnlU

**********
“AI will not wipe out humanity but it's also a risk, says CEO of LLM startup Cohere”

Global AI company Cohere, co-headquartered in Toronto and SF, is one of a crop of AI companies focused on large language models. Aidan Gomez, the company's co-founder, talks to CNBC Senior Tech Correspondent Arjun Kharpal about the technology's development and the applications it will serve.

[ CNBC International TV ]

VIDEO: https://www.youtube.com/watch?v=6LbhUsAwBN0

**********

"Human Extinction: What Are the Risks?"

Correction to what I say at 11 mins 50 seconds: A supervolcano eruption ejects more than 1000 cubic kilometers of matter (not 1000 cubic meters). Sorry about that!

What do we know about the risks of human going extinct? In today's video I collect what we know about the frequency of natural disasters and just how they would kill us, and estimates for man-made disasters.

[ Sabine Hossenfelder ]

VIDEO: https://www.youtube.com/watch?v=nQVgt5eFMh4

**********

"Asilomar AI Principles"

The Asilomar AI Principles, coordinated by FLI and developed at the Beneficial AI 2017 conference, are one of the earliest and most influential sets of AI governance principles.

These principles were developed in conjunction with the 2017 Asilomar conference (videos here), through the process described here.

Click here to see this page in other languages: Chinese German Japanese Korean Russian

Artificial intelligence has already provided beneficial tools that are used every day by people around the world. Its continued development, guided by the following principles, will offer amazing opportunities to help and empower people in the decades and centuries ahead. [more...]

[ Future Of Life ]

WEB PAGE: https://futureoflife.org/open-letter/ai-principles/

**********

“AI-box Experiment”

The AI-box experiment is a thought experiment and roleplaying exercise devised by Eliezer Yudkowsky to show that a suitably advanced artificial intelligence can convince, or perhaps even trick or coerce, people into “releasing” it — that is, allowing it access to infrastructure, manufacturing capabilities, the Internet, and so on. This is one of the points in Yudkowsky's work at creating a friendly artificial intelligence (FAI), so that when “released” an AI won't try to destroy the human race for one reason or another.

Note that despite Yudkowsky's wins being against his own acolytes and his losses being against outsiders, he considers the (unreleased) experimental record to constitute evidence supporting the AI-box hypothesis, rather than evidence as to how robust his ideas seem if you don't already believe them. [more...]

[ Wikipedia ]

WEB PAGE: https://rationalwiki.org/wiki/AI-box_experiment

**********

“10 DISTURBING AI Breakthroughs Coming In 2024”

AI is advancing at a rapid rate that is impossible to keep up with. Come explore some amazing and scary possibilities we may see from AI in 2024.

VIDEO: https://www.youtube.com/watch?v=HE7fcWpvZyg

**********

“Musk claims Neuralink patient doing OK with implant, can move mouse with brain”

Medical ethicists alarmed by Musk being “sole source of information” on patient.

Neuralink co-founder Elon Musk said the first human to be implanted with the company's brain chip is now able to move a mouse cursor just by thinking.

“Progress is good, and the patient seems to have made a full recovery, with no ill effects that we are aware of. Patient is able to move a mouse around the screen

by just thinking," Musk said Monday during an X Spaces event, according to Reuters.

Musk's update came a few weeks after he announced that Neuralink implanted a chip into the human. The previous update was also made on X, the Musk-owned social network formerly named Twitter.

Musk reportedly said during yesterday's chat, "We're trying to get as many button presses as possible from thinking. So that's what we're currently working on is: can you get left mouse, right mouse, mouse down, mouse up... We want to have more than just two buttons."

Neuralink itself doesn't seem to have issued any statement on the patient's progress. We contacted the company today and will update this article if we get a response.

"Basic ethical standards" not met. [more...]

[ JON BRODKIN ]

ARTICLE: https://arstechnica.com/tech-policy/2024/02/musk-claims-neuralink-patient-doing-ok-with-implant-can-move-mouse-with-brain/

**********
"First Neuralink Implanted & Where Other Tech Giants Are Headed w/ | EP #85"

In this episode, Peter and Salim dive into the craziest news in robotics, generative AI robots, Neuralink, and the Uncanny Valley of robots.

Salim Ismail is a serial entrepreneur and technology strategist well known for his expertise in Exponential organizations. He is the Founding Executive Director of Singularity University and the founder and chairman of ExO Works and OpenExO.

[ Peter H. Diamandis ]

INTERVIEW: https://www.youtube.com/watch?v=WO7kzoFSRgY

**********
"Google pauses AI tool Gemini's ability to generate images of people after historical inaccuracies"

Google says it's temporarily suspended the ability of Gemini, its flagship generative AI suite of models, to generate images of people while it works on

updating the technology to improve the historical accuracy of outputs involving depictions of humans.

In a post on the social media platform X, the company announced what it couched as a “pause” on generating images of people — writing that it’s working to address “recent issues” related to historical inaccuracies.

“While we do this, we’re going to pause the image generation of people and will re-release an improved version soon,” it added.

Google launched the Gemini image generation tool earlier this month. However examples of it generating incongruous images of historical people have been finding their way onto social media in recent days — such as images of the U.S. Founding Fathers depicted as American Indian, Black or Asian — leading to criticism and even ridicule.

Writing in a post on LinkedIn, Paris-based venture capitalist Michael Jackson joined the pile-on today — branding Google’s AI as “a nonsensical DEI parody”. (DEI standing for ‘Diversity, Equity and Inclusion.’)

In a post on X yesterday, Google confirmed it was “aware” the AI was producing “inaccuracies in some historical image generation depictions”, adding in a statement: “We’re working to improve these kinds of depictions immediately. Gemini’s AI image generation does generate a wide range of people. And that’s generally a good thing because people around the world use it. But it’s missing the mark here.”

Generative AI tools produce outputs based on training data and other parameters, such as model weights.

Such tools have more often faced criticism for producing outputs that are biased in more stereotypical ways — such as overtly sexualized imagery of women or by responding to prompts for high status job roles with imagery of white men.

An earlier AI image classification tool made by Google caused outrage, back in 2015, when it misclassified Black men as gorillas. The company promised to fix the issue but, as Wired reported a few years later, its ‘fix’ was pure workaround: With Google simply blocking the tech from recognizing gorillas at all.

[ Natasha Lomas ]

**********

“Edward Snowden and Ben Goertzel on the AI Explosion and Data Privacy”

Former defense contractor and activist Snowden and cognitive scientist Goertzel discuss the surveillance implications of recent advancements in AI at Consensus 2023.

[ CoinDesk ]

VIDEO: https://www.youtube.com/watch?v=4H_iWpWG_c0

**********

“Edward Tenner: Unintended consequences”

Every new invention changes the world -- in ways both intentional and unexpected. Historian Edward Tenner tells stories that illustrate the under-appreciated gap between our ability to innovate and our ability to foresee the consequences.

[ TED Talk ]

PRESENTATION: https://www.youtube.com/watch?v=rGaj2VImQec

**********

“Exponential Growth Exposed”

Change Management & Communication at Sidel
Published Mar 13, 2015

Disruptive businesses, technologies and change processes are very often associated with exponential growth. It is not easy to really grasp the concept as we are usually exposed to linear growth (or regression) processes. It is important though to really understand the completely different nature of exponential growth vs. linear, as they have very different implications on how we need to act to deal and take advantage of them.

The following example of exponential growth has been floating around the internet for years and I’ve read it on many blogs and books. It’s an interesting analogy as to how exponential growth works…

“Imagine a magic pipette. It is magic because every drop of water that comes out of it will double in size every minute. So the first minute there is one drop, the second minute there are two drops, the third minute four drops, the fourth minute

eight drops and so on… This is an example of exponential growth. Now, imagine a normal sized football stadium. In this stadium you are sitting on the seat at the very top of the stadium, with the best overview of the whole stadium. To make things more interesting, imagine the stadium is completely water-tight and that you cannot move from your seat. The first drop from the magic pipette is dropped right in the middle of the field, at 12pm. Here's the question: Remembering that this drop grows exponentially by doubling in size every minute, how much time do you have to free yourself from the seat and leave the stadium before the water reaches your seat at the very top? Think about it for a moment. Is it hours, days, weeks, months?

The answer: You have exactly until 12:49pm. It takes this tiny magic drop less than 50 minutes to fill a whole football stadium with water. This is impressive! But it gets better: At what time do you think the football stadium is still 93% empty? Take a guess.

The answer: At 12:45pm. So, you sit and watch the drop growing, and after 45 minutes all you see is the playing field covered with water. And then, within four more minutes, the water fills the whole stadium. This means that you think you are safe because it seems that you have plenty of time left, whereas due to the exponential growth you really have to take immediate action if you want to have any chance of getting out of this situation."

Exponential growth is therefore not only rapid, but we also have very little time in which to react to any situation before it's too late.

Exponential growth thinking is relevant in many fields, from pandemic management to app launch and scale-up. For instance when you are developing an app and have to correctly size cloud resources to serve it, it is worth pausing for a moment to consider what this rapid growth could mean in real terms. To be very successful we need to have a cloud platform that's elastic enough to provide consistent continuity of service to an exponentially growing audience.

[ Giuseppe Geneletti ]

**********

"Exponential Growth Explained"

What happens when an economy reaches its peak? Chis Martenson explains the concept and danger of exponential growth.

[ Chis Martenson ]

VIDEO: https://www.youtube.com/watch?v=6fqYMzFqntg&t=135s

**********

“The Debt Treadmill: How Exponential Debt Growth Spells Trouble”

Let’s roll up our sleeves and dive headfirst into the powerful world of exponentials compounding. You see, when something grows over time, whether it’s population, the demand for good ol’ oil, debt, or even the money supply, and you chart that baby over time, well, it’s got a shape to it, just like a hockey stick. Now, if that growth’s happening on a percentage basis, that’s what we call exponential growth. Let me illustrate this with a little tale inspired by the brilliant Dr. Albert Bartlett.

Imagine I’ve got a magic eye dropper, and I plunk down a single drop of water on your left hand. But here’s the kicker: that drop doubles in size every single minute. First minute, nothing much, just two tiny drops. Keep that doubling up, and in six minutes, you’ve got yourself enough water to fill a dinky thimble.

Now, hold onto your hats, folks, ’cause we’re taking that magical eye dropper to the heart of Yankee Stadium, right on that pitcher’s mound, high noon. Picture this: the park’s sealed up tight, and I’ve shackled you to the nosebleeds. Your mission? Break free before that stadium turns into one colossal water park. How long do you have before you drown? Well, not long at all, my friends, ’cause by 12:50 on the very same day, that itty-bitty drop of water will have turned Yankee Stadium into a ginormous pool. But here’s the real kicker – at 12:45, the stadium’s still 93% empty, and that’s when you’ll realize time isn’t on your side. This, my friends, is the essence of exponential growth: slow and steady at first, but when it hits that vertical climb, it’s a race against the clock. Dr. Bartlett once said, “the greatest shortcoming of the human race is the inability to understand the exponential function.” And he hit the nail on the head. So, remember, once you’re on that vertical ride, time’s your most precious asset.

To provide an easier comprehension of this subject I did the math.  Here is what the horrifying progression looks like:

– At 1 minute: 2 drops of water
– At 2 minutes: 4 drops of water
– At 3 minutes: 8 drops of water
– At 4 minutes: 16 drops of water
– At 5 minutes: 32 drops of water
– At 6 minutes: 64 drops of water
– At 7 minutes: 128 drops of water
– At 8 minutes: 256 drops of water
– At 9 minutes: 512 drops of water

– At 10 minutes: 1,024 drops of water

So, after 10 minutes, we poured in 1,024 drops of water of water.

Let’s continue the example with drops of water for 40 more minutes:

– At 11 minutes: 2,048 drops of water
– At 12 minutes: 4,096 drops of water
– At 13 minutes: 8,192 drops of water
– At 14 minutes: 16,384 drops of water
– At 15 minutes: 32,768 drops of water
– At 16 minutes: 65,536 drops of water
– At 17 minutes: 131,072 drops of water
– At 18 minutes: 262,144 drops of water
– At 19 minutes: 524,288 drops of water
– At 20 minutes: 1,048,576 drops of water
– At 21 minutes: 2,097,152 drops of water
– At 22 minutes: 4,194,304 drops of water
– At 23 minutes: 8,388,608 drops of water
– At 24 minutes: 16,777,216 drops of water
– At 25 minutes: 33,554,432 drops of water
– At 26 minutes: 67,108,864 drops of water
– At 27 minutes: 134,217,728 drops of water
– At 28 minutes: 268,435,456 drops of water
– At 29 minutes: 536,870,912 drops of water
– At 30 minutes: 1,073,741,824 drops of water
– At 31 minutes: 2,147,483,648 drops of water
– At 32 minutes: 4,294,967,296 drops of water
– At 33 minutes: 8,589,934,592 drops of water
– At 34 minutes: 17,179,869,184 drops of water
– At 35 minutes: 34,359,738,368 drops of water
– At 36 minutes: 68,719,476,736 drops of water
– At 37 minutes: 137,438,953,472 drops of water
– At 38 minutes: 274,877,906,944 drops of water
– At 39 minutes: 549,755,813,888 drops of water
– At 40 minutes: 1,099,511,627,776 drops of water
– At 41 minutes: 2,199,023,255,552 drops of water
– At 42 minutes: 4,398,046,511,104 drops of water
– At 43 minutes: 8,796,093,022,208 drops of water
– At 44 minutes: 17,592,186,044,416 drops of water
– At 45 minutes: 35,184,372,088,832 drops of water
– At 46 minutes: 70,368,744,177,664 drops of water
– At 47 minutes: 140,737,488,355,328 drops of water

– At 48 minutes: 281,474,976,710,656 drops of water
– At 49 minutes: 562,949,953,421,312 drops of water
– At 50 minutes: 1,125,899,906,842,624 drops of water

So, after 50 minutes of doubling the amount of water poured into the stadium every minute, we have a whopping 1,125,899,906,842,624 drops of water of water in the stadium. That's more water than you can probably imagine! This example illustrates how exponential growth can lead to incredibly large numbers very quickly.

[more…]

ARTICLE: https://www.vantagepointsoftware.com/blog/exponential-debt-growth/#:~:text=So%2C%20after%2050%20minutes%20of,incredibly%20large%20numbers%20very%20quickly.

**********
"AI Expert Stuart Russell Breaks Down Artificial Intelligence in Films & TV | Break It Down"

Groundbreaking Professor of Computer Science Stuart Russell OBE joins us to break down depictions of AI in movies and TV shows, including Terminator 2, Ex Machina, Jexi, WALL-E and Black Mirror. Order Stuart's AI book Human Compatible here: https://bit.ly/46OrHkS

Humans dream of super-intelligent machines. But what happens if we actually succeed? Creating superior intelligence would be the biggest event in human history. Unfortunately, according to the world's pre-eminent AI expert, it could also be the last.

In this groundbreaking book, Stuart Russell sets out why he has come to consider his own discipline an existential threat to humanity, and how we can change course before it's too late. In brilliant and lucid prose, he explains how AI actually works and its enormous capacity to improve our lives - and why we must never lose control of machines more powerful than we are. Russell contends that we can avert the worst threats by reshaping the foundations of AI to guarantee that machines pursue our objectives, not theirs. Profound, urgent and visionary, Human Compatible is the one book everyone needs to read to understand a future that is coming sooner than we think.

[ Penguin Books UK ]

PRESENTATION: https://www.youtube.com/watch?v=xeYk9T8yOic

**********

“Is AI Just Another Tool, or Something Else?”

We must ask how the technology fits into the Creation story lest it takes away from meaningful work and connection.

It’s not uncommon to hear artificial intelligence described as a new “tool” that extends and expands our technological capabilities. Already there are thousands of ways people are utilizing artificial intelligence. All tools help accomplish a task more easily or efficiently. Some tools, however, have the potential to change the task at a fundamental level.

This is among the challenges presented by AI. If in the end it is not clear what AI is helping us to achieve more efficiently, this emerging technology will be easily abused. AI’s potential impact on education is a prime example.

Since the days of Socrates, the goal of education was not only for students to gain knowledge but also the wisdom and experience to use that knowledge well. Whether the class texts appeared on scrolls or screens mattered little. Learning remained the goal, regardless of the tools used.

In a recent article at The Hill, English professor Mark Massaro described a “wave” of chatbot cheating now making it nearly impossible to grade assignments or to know whether students even complete them. He has received essays written entirely by AI, complete with fake citations and statistics but meticulously formatted to appear legitimate. In addition to hurting the dishonest students who aren’t learning anything, attempts to flag AI-generated assignments, a process often powered by AI, have the potential to yield false positives that bring honest students under suspicion.

Some professors are attempting to make peace with the technology, encouraging students to use AI-generated “scaffolding” to construct their essays. However, this is kind of like legalizing drugs: There’s little evidence it will cut down on abuse.

Consider also the recent flood of fake news produced by AI. In an article in The Washington Post, Pranshu Verma reported that “since May, websites hosting AI-created false articles have increased by more than 1,000 percent.” According to one AI researcher, “Some of these sites are generating hundreds if not thousands of articles a day… This is why we call it the next great misinformation superspreader.”

Sometimes, this faux journalism appears among otherwise legitimate articles. Often, the technology is used by publications to cut corners and feed the content machine. However, it can have sinister consequences.

A recent AI-generated story alleged that Israeli prime minister Benjamin Netanyahu's psychiatrist had committed suicide. The fact that this psychiatrist never existed didn't stop the story from circulating on TV, news sites, and social media in several languages. When confronted, the owners of the site said they republished a story that was "satire," but the incident demonstrates that the volume of this kind of fake content would be nearly impossible to police.

Of course, there's no sense in trying to put the AI genie back in a bottle. For better or worse, the technology is here to stay. We must develop an ability to evaluate its legitimate uses from its illegitimate uses. In other words, we must know what AI is for, before experimenting with what it can do.

That will require first knowing what human beings are for. For example, Genesis is clear (and research confirms) that human beings were made to work. After the fall, toil "by the sweat of your brow" is a part of work. The best human inventions throughout history are the tools that reduce needless toil, blunt the effects of the curse, and restore some dignity to those who work.

We should ask whether a given application of AI helps achieve worthy human goals—for instance, teaching students or accurately reporting news—or if it offers shady shortcuts and clickbait instead. Does it restore dignity to human work, or will it leave us like the squashy passengers of the ship in Pixar's Wall-E—coddled, fed, entertained, and utterly useless?

Perhaps most importantly, we must govern what AI is doing to our relationships. Already, our most impressive human inventions—such as the printing press, the telephone, and the internet—facilitated more rapid and accurate human communication, but they also left us more isolated and disconnected from those closest to us. Obviously, artificial intelligence carries an even greater capacity to replace human communication and relationships (for example, chatbots and AI girlfriends).

In a sense, the most important questions as we enter the age of AI are not new. We must ask, what are humans for? And, how can we love one another well? These questions won't easily untangle every ethical dilemma, but they can help distinguish between tools designed to fulfill the creation mandate and technologies designed to rewrite it.

[ John Stonestreet and Shane Morris ]

**********

“The Eye: A Masterpiece of God”

Eyes have been called the “windows to the soul” because we use them to express emotion and they are quite beautiful. These are secondary functions but still very important. Our eye’s most important function, of course, is giving us vision. So how does this work?

We have two lenses, called the cornea and lens proper, which work much like lenses on a camera, only much better. These lenses form in the womb when embryonic skin turns into a clear window. The normal blood vessels, hair, and sweat glands that grow in your skin are missing from this small area, though it’s filled with sensitive nerves. These lenses change shape depending on the available light and can automatically adjust over a range of 10 billion to one!

Twelve muscles, six for each eye, control the movement of our eyes and lenses. If these muscles or our eyes were even slightly misaligned we would see double. It’s as if our eyes have been designed with incredible precision!

The cornea catches the light bouncing off objects and it travels to the pupil, a round hole in our eye. The iris, the colored part of your eye, regulates the amount of light allowed into the pupil by opening and closing. Now, this changes the size of the pupil.

The light then goes to the lens proper which is flexible and focuses the light on our retina, a thin layer of tissue with millions of light-sensing nerve cells called rods and cones. These cells convert light into electrical impulses which travel up the optic nerve to be understood by the brain.

Whew! That’s quite the process and all that happens without you having to “bat an eyelash,” so to speak!

Now, some proponents of Darwinian evolution will argue that the human eye is actually an inefficient design despite its superiority to any man-made camera, its self-cleaning and self-repairing equipment, and its precision. The argument goes this way: we have an inverted retina so the light has to pass through several layers before it gets to the photoreceptors. But if our eyes had been designed without these extra layers, our retinas would be more easily damaged by bright light and heat. Sure, we might be able to see in the dark better, but daytime sunlight might be almost painful, so, the design we have is actually the best one! I believe that we were created with purpose, and that the God of the Bible thought

of everything! And if you want to see better in the dark, don't look to evolution to help you adapt over the next million years - That's what carrots are for.

[ David Rives ]

PRESENTATION: https://www.youtube.com/watch?v=7HY_v9O7Bpc

**********

"The "Eyes" Have It!"

"How could anyone possibly believe that something so obviously designed for a specific purpose is a product of evolution?"

This rather rhetorical question (or so I thought) was posed by my daughter after she shared scads of information from her Anatomy & Physiology book, enthusiastically declaring how absolutely amazing the intricacies of the human eye are. "Cool!" I thought to myself, "but what do I really know about it" (certainly, not as much as I wanted to)? So, being a lover of learning, I had to start investigating this intriguing topic--the human eye.

Though God doesn't talk specifically about the spleen, gallbladder, or the thyroid, when He does talk about anatomy in His Word, it shows us that He takes a special interest in us and how He designed us.

As a believer in Christ, I am interested first of all in what God's Word has to say on the subject, and because of my biblical worldview, I believe that the eye is one of the most miraculous products of God's amazing design within the human body. [more...]

[ Reasons For Hope ]

ARTICLE: https://www.rforh.com/blog/2021/10/08/the-eyes-have-it

**********

"The Human Eye; a masterpiece of design"

Overview
The human eye is extremely complex—a perfect and interrelated system of about 40 individual subsystems, including the retina, pupil, iris, cornea, lens and optic nerve. For instance, the retina has approximately 137 million special cells that respond to light and send messages to the brain. About 130 million of these cells look like rods and handle the black and white vision. The other seven million are cone-shaped and allow us to see in color. The retina cells receive light

impressions, which are converted to electric pulses and sent to the brain via the optic nerve. A special section of the brain, called the visual cortex transforms the pulses to color, contrast, depth, etc., which allows us to see 'pictures' of our world in three dimensions. Amazingly, the eye, the optic nerve and the brain's visual cortex are totally separate and distinct subsystems. Yet, together they capture, deliver and interpret up to 1.5 million pulse messages every milli-second (one thousandth of a second)! [more...]

[ A Defence Of The Bible ]

ARTICLE: https://www.adefenceofthebible.com/2017/08/15/the-human-eye-a-masterpiece-of-design/

**********
"God's Greatest Creation"

You saw me before I was born. Every day of my life was recorded in your book. Every moment was laid out before a single day had passed.
—Psalm 139:16

Without question, people are God's greatest creation. We are His crowning achievement. In fact, the psalmist David wrote about the intricacies of the human body that God created. He said, "You made all the delicate, inner parts of my body and knit me together in my mother's womb. Thank you for making me so wonderfully complex! Your workmanship is marvelous—how well I know it" (Psalm 139:13–14 NLT).

David continued, "You watched me as I was being formed in utter seclusion, as I was woven together in the dark of the womb. You saw me before I was born. Every day of my life was recorded in your book. Every moment was laid out before a single day had passed" (verses 15–16 NLT).

As we look at Scripture, it appears that God has a plan for each of us, even before we were conceived. The prophet Jeremiah wrote, "The Lord gave me this message: 'I knew you before I formed you in your mother's womb. Before you were born I set you apart and appointed you as my prophet to the nations'" (Jeremiah 1:4–5 NLT).

These verses, among others, certainly lay to rest any warped concept that the Bible would somehow allow for abortion.

This masterpiece of God's creation, the human being, is incredible. Scientists estimate that adult human bodies contain 16 trillion cells, all carefully organized to perform life's various functions in harmony.

Consider these statistics about the human body and the amazing things it is capable of: The nose can recognize 10,000 different aromas. The tongue has about 6,000 taste buds. And the brain contains 10 billion nerve cells. Each nerve cell connects to as many as 10,000 other nerve cells throughout the body. In fact, the body has so many blood vessels that the combined length could circle the planet two and a half times.

God has custom-designed each of us with our own DNA blueprint, which every cell contains. And if you were to write out your individual blueprint in a book, it would require an estimated 200,000 pages. God, of course, knows about every word on every page.

We have the astounding capacity to store millions of bits of information, keep it in order, and recall it when necessary. We are "wonderfully complex," as Psalm 139 tells us. And God's plan for those who put their faith in Christ is even more amazing. We don't have to be afraid because His motive is always love for us.

Ephesians 2:6–7 tells us, "He raised us from the dead along with Christ and seated us with him in the heavenly realms because we are united with Christ Jesus. So, God can point to us in all future ages as examples of the incredible wealth of his grace and kindness toward us, as shown in all he has done for us who are united with Christ Jesus" (NLT).

God wants to spend all of eternity revealing to us His kindness, goodness, and grace. He wants to spend eternity showing us how much He loves us.

[ Greg Laurie ]

**********

"Is the Bible True?"
Evidence for the bible or proof the bible is true, is everywhere, the questions is, when someone asks, is the Bible true or, is the Bible real, do you know how to answer? Just asking someone to trust the Bible, may not be enough, they want evidence for the Bible. No problem. There are many ways to answer, is the Bible reliable, but in this video Pastor Nelson looks at the story of David and Bathsheba in 2 Samuel 11 as evidence for the Bible that helps answer the question, is the Bible True.

[ Bible Munch ]

Video: https://www.youtube.com/watch?v=9iyMjG0Ru1E

**********

“Is the Bible trustworthy?”

[ Todd Friel ]

Video: https://www.youtube.com/watch?v=jCmGlUCgEX8

**********

“Is The Bible Trustworthy?”

How do we know that the Bible is reliable and that the words contained in today’s Bibles convey the same message as the original documents?

[ Ready 4 Eternity ]

Video: https://www.youtube.com/watch?v=20yiTTJVaqg

**********

“Is the New Testament Reliable?”

For more information, read Evidence that Demands A Verdict, co-written by Sean and Josh McDowell

[ SeanMcDowell ]

Video: https://www.youtube.com/watch?v=bs4s2i_bWEw

_____

_____

## SONGS

_____

_____

“**Masterpiece**”

Heartbreaks a bittersweet sound
Know it well

It's ringing in my ears
And I can't understand
Why I'm not fixed by now
Begged and I pleaded
Take this pain but I'm still bleeding

Heart trusts you for certain
Head says it's not working
I'm stuck here still hurting
But you tell me

You're making a masterpiece
You shaping the soul in me
You're moving where I can't see
And all I am is in your hands
You're taking me all apart
Like it was your plan from the start
To finish your work of art for all to see you're making a masterpiece

Guess I'm your canvas
Beautiful black and blue
Painted in mercy's hue
I don't see past this
But you see me now
Who I'll be then
There at the end
Standing there as

Your Masterpiece
You're shaping the soul in me
You're moving where I can't see
And all I am is in your hands
You're taking me all apart
Like it was your plan from the start
To finish your work of art for all to see
You're making a masterpiece
You're making a masterpiece

Even though I'm hurting
I'll let you keep working
You're making a masterpiece
You're shaping the soul in me
You're moving where I can't see

And all I am is in your hands

You're taking me all apart
Like it was your plan from the start
To finish your work of art for all to see you're making a masterpiece
You're making a masterpiece

I'll be your masterpiece

[ Danny Gokey - “Rise” album ]

SONG: https://www.youtube.com/watch?v=culDfUqiKDE

**********

“**Masterpiece**”

You knew me from the start
I came from Your heart
Breathed Your breath in me
Took me from the clay
Formed me in Your shape
Said that I was wonderfully made

You said come alive
You said come alive
Living soul, breathe breath of life

I will love You
I will love You
I will love You
I can’t stop loving you
I will love You
I will love You
I will love You
I can’t stop loving you

You knew me from the start
I came from Your heart
Breathed Your breath in me
Took me from the clay
Formed me in Your shape
Said that I was wonderfully made

I am Your masterpiece
You made no mistake with me
I am Your masterpiece
You made no mistake with me
I am Your masterpiece
You made no mistake with me
I am Your masterpiece
You made no mistake with me

[ Destiny Worship Music - “Just Jesus” album ]

SONG: https://www.youtube.com/watch?v=Fbpm_UxsCdc

**********

“**Poem Of Your Life**”

Life is a song we must sing with our days
A poem with meaning more than words can say
A painting with colors no rainbow can tell
A lyric that rhymes either heaven or hell!

We are living letters that doubt desecrates
We’re the notes of the song of the chorus of faith
God shapes every second of our little lives
And minds every minute as the universe waits by

The pain and the longing
The joy and the moments of light
Are the rhythm and rhyme
The free verse of the poem of life

So look in the mirror and pray for the grace
To tear off the mask, see the art of your face
Open your ear lids to hear the sweet song
Of each moment that passes and pray to prolong

Your time in the ball of the dance of your days
Your canvas of colors of moments ablaze
With all that is holy, with the joy and the strife
With the rhythm and rhyme of the poem of your life
With the rhythm and rhyme of the poem of your life

The pain and the longing

The joy and the moments of light
Are the rhythm and rhyme
The free verse of the poem of life.

[ Michael Card - “Poiema” album ]

SONG: https://www.youtube.com/watch?v=-Sjb2AhGoi0

**********

“**Human Nature**”

We're not God's problem, we are God's children
We are not hopeless, we're what He's building
We can be unfinished and still be His image
To bring in the light to the land of the living

We're not God's problem, we are His children
So don't let nobody tell you otherwise
The truth is in your Father's eyes
And anything else is a lie
It's a lie, it's a lie, it's a lie

When the Savior meets human nature (human nature)
Somehow the world don't seem so empty any more
When the Savior meets human nature (human nature)
That's all we gotta do is open the door
And say, come on in, come on in, come on in
And say, come on in, come on in, come on in

We're not God's problem, we are God's family
We're at the table, like it's thanksgiving
He's not mad, but He loves you madly
Can see what you've done, but He can see what you will be
We're not God's problem, we are God's family

When the Savior meets human nature (human nature)
Somehow the world don't seem so empty any more
When the Savior meets human nature (human nature)
That's all we gotta do is open the door
And say, come on in, come on in, come on in
And say, come on in, come on in, come on in

Here's the thing

Being human means making mistakes
But don't be mistaken
Being human means you're worthy of love, you're worthy of love

So don't let nobody tell you otherwise
The truth is in your Father's eyes
And anything else is a lie
It's a lie, it's a lie, it's a lie

When the Savior meets human nature (human nature)
Somehow the world don't seem so empty any more
When the Savior meets human nature (human nature)
That's all we gotta do is open the door
And say, come on in, come on in, come on in (Savior)
And say, come on in, come on in, come on in (human nature)
And say, come on in, come on in, come on in (Savior)
He's out there knocking, so let's open the door

[ Brandon Heath - “Enough Already” album ]

SONG: https://www.youtube.com/watch?v=35JU_EwcUhQ

**********

“**O To Be Like Thee**”

O to be like Thee, blessed Redeemer,
This is my constant longing and prayer;
Gladly I’ll forfeit all of earth’s treasures
Jesus, Thy perfect likeness to wear.

CHORUS
O to be like Thee, O to be like Thee,
Blessed Redeemer, pure as Thou art.
Come in Thy sweetness, come in Thy fullness;
Stamp Thine own image deep on my heart.

O to be like Thee, full of compassion,
Loving, forgiving, tender and kind,
Helping the helpless, cheering the fainting,
Seeking the wand’ring sinner to find.

O to be like Thee, lowly in spirit,
Holy and harmless, patient and brave;

Meekly enduring cruel reproaches,
Willing to suffer, others to save.

[ Nathan Drake - “Reawaken Hymns” album ]

SONG: https://www.youtube.com/watch?v=V6uJtUD43c0

---

## APOLOGETIX SONGS

---

“**Something Of Value**”
(Parody of “Something About You” by Level 42)

Ooh-ooh, ooh-ooh, ooh, ooh, ooh, ooh
Ooh-ooh, ooh-ooh, ooh, ooh, ooh, ooh

How proud can we be when our lives start out with parents
Passing down traits — with some good, some not
And our ways stray much too often?
But taking His grace is the start of finding redemption
We’re all made in the image of God — like the humans at the fall

Oh, drawn into a scheme — we run to substitutions
Gold, diamond rings — and fancy kinds of jewels
But there is something of value greater — your life
I wouldn’t preach without proof, baby — would I?

Since that first garden was concealed, no one can say that
We didn’t deal with myriad sins or deserved the Tree of Life
Gone — Adam and Eve could have stayed in there forever
If not for the stuff — (they) did that’s so wrong — the only humans at the fall

Though it’s strange and weird, we have their constitution
Romans you need to hear — in 5 verse 12, it’s true
But there is something of value they had — no lie
(Something about — the way you are designed)
Those ones who sinned defiled you — made you to die

(Though I can see — we're bound to sin and die)
But God shared something of value — later — so high
(Something about — the grace that God provides)
Come like a little child, too — baby — to Christ — yeah-eh, yeah
(I couldn't list — the value; it's too high)

Ooh-ooh, ooh-ooh, ooh, ooh, ooh, ooh
Ooh-ooh, ooh-ooh, ooh, ooh, ooh, ooh
Come now to Psalms if doubt proof, yeah-eh, yeah
Better look at 51:2 through 5

[ ApologetiX – "Never Before, but Then Again…" album ]

SONG: https://www.youtube.com/watch?v=3MrI80PgiQ8

**********

"**You Made Me, So There You Have Me**"
(Parody of "You've Made Me So Very Happy" performed by Blood, Sweat & Tears)

I lost my nerve before -- got scared and no support
But You said, child, there's just one Lord
I know You are the One -- God in Heaven's Only Son
You created me, so Christ -- I'm about to lose my rights

You made me, so there You have me
I'm so blessed You came into my life

The others searched for Truth -- but when it came they never knew
Pride's been my whole life's pursuit
But You came and You took control -- You blessed my weary soul
You called and told me that -- up with You is where it's at

You made me, so there You have me
I'm so bad You came into my life

Then You saved me -- ay-eh-yeah-eh

I love You so much it seems -- the people think I'm extreme
I can hear -- names I can hear them callin' me
I'll show them love with truth -- I'll lead everyone to You that
Thinks I'm crazy -- thinks I'm crazy

You made me, so there You have me
I'm so bad You came into my life

You made me, so there You have me
You made me, so -- show where You'd have me daily
I'm so blessed You came -- into my life

Hmmm mmm mmm … I wanna thank You, Lord
In every phase of my life … I wanna thank You
You made me -- show where You'd have me
'Cause I wanna spend my life thankin' You
Thank You daily, thank you greatly
Thank You plainly, thank you crazily

[ ApologetiX – "Xit Ego Lopa" album ]

SONG: https://www.youtube.com/watch?v=DctFtlI8xHk

**********

"**Romans 1:20**"
(Parody of "Running on Empty" performed and written by Jackson Browne)

Lookee now back in Romans chapter one, if you will
Look at that and it's clear God provided so many signs, I feel
If you find Romans 1:18 then you're settin' up what I want
I don't know where you're readin' now but just come along

Come along
Romans 1:20 - come along
Come and find
Come along
Romans 1:21 - come and study the lines

God gave proof unto Man yet the people closed their eyes
Science often confuses the truth with some beautiful lies
They've hypnotized almost everyone and they call us rogues, I know
I don't know why they won't turn onto the Romans Road

Come along
Romans 1:20 - come along
Come and find
Come along
Romans 1:21 - come and study the lines

LEAD

Everyone does know - and the Word does show
People seek some reason to believe
I don't know just how anyone can't see
If they ain't got sight - that would seem all right
But other senses besides inform our beliefs

Lookee now, back in Romans chapter one, if you will
I do not have to tell you all just how atheists have appeal
Lookin' round for the answers I could determine who told me the truth
Lookin' into their writings I see there's nothing new

Come along
Romans 1:20 - come along
Come and find
Come along
Romans 1:21 - come and study the lines

What they do's really tempting, you know, they make it look so fine
They love to sniff around, but they're dumb and they're blind
(Come along)
You know, I don't even know how to open their minds
Come and find
Romans 1:21 - come and study the lines

[ ApologetiX – “Conspiracy No. 56” album ]

SONG: https://www.youtube.com/watch?v=688HXCzvNtg

**********

“**Back Talk**”
(Parody of “Black Dog” by Led Zeppelin)

Hey, hey, Moses said that wayward dudes
Posed a major threat for the ancient Jews

Ah, watch out, when they change out things
From His sacred Word, something they do stinks

Hey, hey, baby, when they talk that way
Watch your money clip and keep away

Oh yeah, oh yeah-eh, oh, oh, oh, oh
Oh yeah, oh yeah-eh, oh, oh, oh, oh

God's law is old -- it stands still
But the faint of heart can't accept God's will

I'm surprised -- they turned and fled
Please review how you are led

Ah ah
Ah ah ah ah
Ah ah ah ah
Ah ah ahhhh

Hey, baby, whoa, baby, listen, baby
Modern days are provin' how
Hey, baby, oh, baby, listen, baby
Moses wasn't foolin' now

Listen to this song -- it might sound loud
But the people I mean will try and drown it out

Sin's not funny -- sticks like tar
Start to dwell in it then It's gonna leave a scar

I don't know, but I've been told
A big bad wolf, it can gobble your soul

Oh yeah, oh yeah-eh, oh, oh, oh, oh
Oh yeah, oh yeah-eh, oh, oh, oh, oh

All I ask for -- all I pray
Instead of lovin' the world I wanna love God's Way

Need a compass to hold in my hand
Cause celluloid lies gave me a savage land

Ah ah
Ah ah ah ah
Ah ah ah ah
Ah ah ahhh

Ahhhhh ahhh yeah yeah

[ ApologetiX – “Prehysterical” album ]

SONG: https://www.youtube.com/watch?v=z92ez4bq8Nc

**********

“**Child of God**”
(Parody of “Shining Star” performed by Earth, Wind & Fire)

Yeah, yeah – hey – huh
Bet you wish you were a star
You dream of fame and fancy cars, yeah
But when you're with the Nazarene
Life ain't always such a dream, oh yeah
What you'll be now's not so clear, hey
Yet to Christ you're very dear, yeah

You're a child of God – don't matter who you're not
We're the bride to be – of Jesus, you and me – His church in unity
LEAD

Child, what's gotten into you?
Child it's not quite "what" -- it's "who"?  Yeah!
It's His Spirit there along, yeah
Yeah, makes His Body quick and strong, yeah
We're goin' to mansions of the Son -- yeah
Yeah, sowin' God's Word to everyone
Yeah, God will help, His Spirit will move, yeah
Well, yes, He will, I got my proof, oh yeah, oh yeah
So edify yourself and read
I know you're facin' some adversity
Jesus Christ is greater than
Are you His? Say, yes, I am

You're a child of God – don't matter who you're not
We're the bride to be – of Jesus, you and me
You're a child of God – don't matter who you're not
We're the bride to be – of Jesus, you and me
You're a child of God – don't matter who you're not
We're the bride to be – of Jesus, you and me

Child of God, though you can't see what you're like and soon will be
Child of God, though you can't see what you're like and soon will be

Find First John 3:2 and read --- and Philippians 2:15

[ ApologetiX – “Very Vicarious” album ]

SONG: https://www.youtube.com/watch?v=ChjTfuFEUGQ

**********

“**Very Last City**”
(Parody of “Paradise City” by Guns N’ Roses)

Take me now to the very last city
Where they have safe streets and they’re gold and pretty
(Take me home!) Oh, won’t you please take me home
Take me now to the very last city
Where the massive gates have a pearly finish
(Take me home!) Oh, won’t you please take me home

Brand-new Earth and Heaven under His feet
C’mon! Our King says come to Me
I’m preparin’ a place for all the ones who believe
I’ll see you on the other side
It’s waitin’ at the end of the line

The righteous risen Lord showed the way
You gotta – keep pushin’ toward that fortress of faith
You know it’s – it’s all in Revelation – just you wait
You read in all the chapters you’ll find
21 and 22 are sublime

Take me now to the very last city
Where they have safe streets and they’re gold and pretty
Oh, won’t you please take me home, yeah, yeah
Take me now to the very last city
Where the massive gates have a pearly finish
Take – me – home!

After we’re there in the city, past danger
He’ll wipe our tears, so just try to remember
A certain Gentleman from Nazareth decreed
We have another city naked eyes can’t see
Tell me who you’re gonna believe

Take me now to the very last city

Where they have safe streets and they're gold and pretty
Take — me – home, yeah, yeah
Take me now to the very last city
Where the massive gates have a pearly finish
Oh, won't you please take me home – oh yeah!

No more delay — no more decay
No more dismay — still more to say

Have to remember this important part
Now we can only enter with an open heart
He said, don't wait around for any action to start
You just believe and you'll find
What are you, blind?
He said it all a million times!

Take me now to the very last city
Where they have safe streets and they're gold and pretty
Take — me – home, yeah, yeah
Take me now to the very last city
Where the massive gates have a pearly finish
Oh, won't you please take me home

Take me now to the very last city
Where they have safe streets and they're gold and pretty
Take — me – home, yeah, yeah
Take me now to the very last city
Where the massive gates have a pearly finish
Oh, won't you please take me home
Home!
LEAD

Ohhhhhhh!
I wanna go, I'm gonna go
Oh, won't you please take me home
I wanna see, I wanna be
Oh, won't you please take me home
Take me now to the very last city
Where they have safe streets and they're gold and pretty
Take – me — home!
Take me now to the very last city
Where the massive gates have a pearly finish
Oh, won't you please take me home

Take me now, take me now
Oh, won't you please take me home
I wanna see, I'm gonna be
Oh, won't you please take me home
MINI LEAD

I wanna see — where I'm gonna be
Oh — oh, take me home
Take me now to the very last city
Where they have safe streets and they're gold and pretty
Oh, won't you please take me home
I wanna go, I — wanna go
Oh, won't you please take me home
Yeah, baby! Whee!

[ ApologetiX – "Nichey" album ]

SONG: https://www.youtube.com/watch?v=7XaoPDNDQco

**********

"**Good News/Bad News**"

This is a Gospel presentation and personal testimony of J. Jackson, lead vocalist of ApologetiX from their 20th-anniversary concert. It is available on 20:20 Vision.

VIDEO (audio only): https://www.youtube.com/watch?v=q21Jnaq-EL8

---

---

# DEEP THOUGHTS

---

---

*"Our approach to existential risks cannot be one of trial-and-error. There is no opportunity to learn from errors. The reactive approach-see what happens, limit damages, and learn from experience—is unworkable."*
*[ Nick Bostrom, faculty of Philosophy, Oxford University ]*

*"The AI does not hate you, nor does it love you, but you are made out of atoms which it can use for something else."*

*[ Eliezer Yudkowsky, Research Fellow, Machine Intelligence Research Institute ]*

*“AGI is intrinsically very, very dangerous. And this problem is not terribly difficult to understand. You don't need to be super smart or super well informed, or even super intellectually honest to understand this problem.”*
*[ Michael Vassar, President, Machine Intelligence Research Institute ]*

*“I definitely think that people should try to develop Artificial General Intelligence with all due care. In this case, all due care means much more scrupulous caution than would be necessary for dealing with Ebola or plutonium.”*
*[ Michael Vassar ]*

*“Existential Risk: One where an adverse outcome would either annihilate earth originating intelligent life or permanently and drastically curtail its potential.”*
*[ Nick Bostrom ]*

*“An existential risk is one where humankind as a whole is imperiled. Existential disasters have major adverse consequences for the course of human civilization for all time to come.”*
*[ Nick Bostrom ]*

*“With the possible exception of nanotechnology being released upon the world there is nothing in the whole catalogue of disasters that is comparable to AGI.”*
*[ Eliezer Yudkowsky, Research Fellow, Machine Intelligence Research Institute ]*

*“We are beginning to depend on computers to help us evolve new computers that let us produce things of much greater complexity. Yet we don’t quite understand the process—it's getting ahead of us. We're now using programs to make much faster computers so the process can run much faster. That's what's so confusing- technologies are feeding back on themselves; we're taking off. We're at that point analogous to when single-celled organisms were turning into multi-celled organisms. We are amoebas and we can’t figure out what the hell this thing is that we’re creating.”*
*[ Danny Hillis, founder of Thinking Machines, Inc. ]*

*“If a system has awareness of itself and can create a better version of itself, that's great,” Omohundro told me. “It’ll be better at making better versions of itself than human programmers could. On the other hand, after a lot of iterations, what does it become? I don’t think most AI researchers thought there'd be any danger in creating, say, a chess-playing robot. But my analysis shows that we should think carefully about what values we put in or we'll get something more along the lines of a psychopathic, egoistic, self-oriented entity.”*
*[ Steve Omohundro - AI pioneer and Stanford professor ]*

*“We won't really be able to understand why a superintelligent machine is making the decisions it is making. How can you reason, how can you bargain, how can you understand how that machine is thinking when its thinking in dimensions you can't conceive of?”*
*[ Kevin Warwick, professor of Cybernetics, University of Reading ]*

*“From the standpoint of existential risk, one of the most critical points about Artificial Intelligence is that an Artificial Intelligence might increase in intelligence extremely fast. The obvious reason to suspect this possibility is recursive self-improvement. (Good 1965.) The AI becomes smarter, including becoming smarter at the task of writing the internal cognitive functions of an AI, so the AI can rewrite its existing cognitive functions to work even better, which makes the AI still smarter, including smarter at the task of rewriting itself, so that it makes yet more improvements... The key implication for our purposes is that an AI might make a huge jump in intelligence after reaching some threshold of criticality.”*
*[ Eliezer Yudkowsky, research fellow, Machine Intelligence Research Institute ]*

*“But if the technological Singularity can happen, it will. Even if all the governments of the world were to understand the “threat” and be in deadly fear of it, progress toward the goal would continue. In fact, the competitive advantage—economic, military, even artistic—of every advance in automation is so compelling that passing laws, or having customs, that forbid such things merely assures that someone else will.”*
*[ Vernor Vinge - “The Coming Technological Singularity” ]*

*“What, then, is the Singularity? It’s a future period during which the pace of technological change will be so rapid, its impact so deep, that human life will be irreversibly transformed. Although neither utopian nor dystopian, this epoch will transform the concepts that we rely on to give meaning to our lives, from our business models to the cycle of human life, including death itself.”*
*[ Ray Kurzweil - “The Singularity is Near: When Humans Transcend Biology” ]*

*“In contrast with our intellect, computers double their performance every eighteen months. So, the danger is real that they could develop intelligence and take over the world.”*
*[ Stephen Hawking, Physicist ]*

*“Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended. Is such progress avoidable? If not to be avoided, can events be guided so that we may survive?”*
*[ Vernor Vinge, author, professor, computer scientist ]*

*“Day by day, however, the machines are gaining ground upon us; day by day we are becoming more subservient to them; more men are daily bound down as slaves to tend them, more men are daily devoting the energies of their whole lives to the development of mechanical life. The upshot is simply a question of time, but that the time will come when the machines will hold the real supremacy over the world and its inhabitants is what no person of a truly philosophic mind can for a moment question.”*
*[ Samuel Butler, nineteenth-century English poet and author ]*

*“How can we be so confident that we will build super-intelligent machines? Because the progress of neuroscience makes it clear that our wonderful minds have a physical basis, and we should have learned by now that our technology can do anything that's physically possible. IBM's Watson, playing Jeopardy as skillfully as human champions, is a significant milestone and illustrates the progress of machine language processing. Watson learned language by statistical analysis of the huge amounts of text available on-line. When machines become powerful enough to extend that statistical analysis to correlate language with sensory data, you will lose a debate with them if you argue that they don't understand language.”*
*[ Bill Hibbard, AI scientist ]*

*“Is it really so far-fetched to believe that we will eventually uncover the principles that make intelligence work and implement them in a machine, just like we have reverse engineered our own versions of the particularly useful features of natural objects, like horses and spinnerets? News flash: the human brain is a natural object.”*
*[ Michael Anissimov, MIRI Media Director ]*

*“Despite all it has borrowed from Christianity, transhumanism is ultimately fatalistic about the future of humanity. Its rather depressing gospel message insists that we are inevitably going to be superseded by machines and that the only way we can survive the Singularity is to become machines ourselves.”*
*[ Meghan O Gieblyn ]*

*“Both because of its superior planning ability and because of the technologies it could develop, it is plausible to suppose that the first superintelligence would be very powerful. Quite possibly, it would be unrivalled: it would be able to bring about almost any possible outcome and to thwart any attempt to prevent the implementation of its top goal. It could kill off all other agents, persuade them to change their behavior, or block their attempts at inter-ference. Even a “fettered superintelligence” that was running on an isolated computer, able to interact with the rest of the world only via text interface, might be able to break out of its*

*confinement by persuading its handlers to release it. There is even some preliminary experimental evidence that this would be the case."*
*[ Nick Bostrom, Future of Humanity Institute, Oxford University ]*

*"I'm going to predict that we are just a few years away from a major catastrophe being caused by an autonomous computer system making a decision."*
*[ Wendall Wallach, ethicist, Yale University ]*

*"There are no catastrophes that loom before us which cannot be avoided; there is nothing that threatens us with imminent destruction in such a fashion that we are helpless to do something about it. If we behave rationally and humanely: if we concentrate coolly on the problems that face all of humanity, rather than emotionally on such nineteenth century matters as national security and local pride; if we recognize that it is not one's neighbors who are the enemy, but misery, ignorance, and the cold indifference of natural law-then we can solve all the problems that face us. We can deliberately choose to have no catastrophes at all."*
*[ Isaac Asimov ]*

*"The Dark Ages may return, the Stone Age may return on the gleaming wings of Science, and what might now shower immeasurable material blessings upon mankind, may even bring about its total destruction. Beware, I say; time may be short. Do not let us take the course of allowing events to drift along until it is too late."*
*[ Winston Churchill ]*

*"A new type of thinking is essential if mankind is to survive and move toward higher levels."*
*[ Albert Einstein ]*

*"More than any other time in history mankind faces a cross-roads. One path leads to despair and utter hopelessness, the other to total extinction. Let us pray we have the wisdom to choose correctly."*
*[ Woody Allen ]*

*"When you learn how to love yourself, you will be better able to love others as yourself."*
*[ Mark Besh ]*

*"Life is eternity's sunrise."*
*[ Mark Besh ]*

---

**RELATED SCRIPTURE VERSES**:

**Humanity**:
https://www.openbible.info/topics/humanity

**Extinction**:
https://www.openbible.info/topics/extinction

**Risks**:
https://www.openbible.info/topics/risks

**Reducing Risks**:
https://www.openbible.info/topics/reducing_risks

**Safety**:
https://www.openbible.info/topics/safety

**Catastrophe**:
https://www.openbible.info/topics/catastrophe

**Existential Threat**:
https://www.openbible.info/topics/existential_threat

**Existential Risk**:
https://www.openbible.info/topics/existential_risk

**Transhumanism**:
https://www.openbible.info/topics/transhumanism

**Immortality**:
https://www.openbible.info/topics/immortality

**Eternal**:
https://www.openbible.info/topics/eternal

**Eternity**:
https://www.openbible.info/topics/eternity

**********
**********

“A quick summary of the Christian “Gospel”:
JESUS’ PROPITIATION made our SINS FORGIVEN and IMPUTED RIGHTEOUSNESS to us so that we have GOD’S ACCEPTANCE into His Heaven and receive ETERNAL LIFE.”
[ Mark Besh ]

**********
**********

Hope you enjoyed some of these insights—share them with your friends and colleagues—so we can have a larger ’pool’ to receive from, and more to share with! Also, remember to include your name as the “source,” if some of this wisdom is of your doing. I would like to give credit where credit is due!

_____
_____

## FOCUS VERSES

_____
_____

*“In the beginning, God created the heavens and the earth.”*
*[ Genesis 1:1 ]*

*“God created man in His own image, in the image of God he created him; male and female he created them. And God blessed them. And God said to them, ‘Be fruitful and multiply and fill the earth and subdue it, and have dominion over the fish of the sea and over the birds of the heavens and over every living thing that moves on the earth.’”*
*[ Genesis 1:27-28) ]*

*“‘The man has now become like one of us, knowing good and evil. He must not be allowed to reach out his hand and take also from the tree of life and eat, and live forever.’ So the Lord God banished him from the Garden of Eden to work the ground from which he had been taken. After he drove the man out, he placed on the east side of the Garden of Eden cherubim and a flaming sword flashing back and forth to guard the way to the tree of life.”*

*[ Genesis 3:22-24 ]*

*"The LORD is God; besides Him there is no other."*
*[ Deuteronomy 4:35 ]*

*"Only fear the LORD and serve him faithfully with all your heart. For consider what great things he has done for you."*
*[ 1 Samuel 12:24 ]*

*"I know that my Redeemer lives… and I myself will see Him with my own eyes."*
*[ Job 19:25-27 ]*

*"What is man that you are mindful of him, and the son of man that you care for him?"*
*[ Psalm 8:3-4 ]*

*"Yet you have made him a little lower than the heavenly beings and crowned him with glory and honor. You have given him dominion over the works of your hands; you have put all things under his feet, all sheep and oxen, and also the beasts of the field, the birds of the heavens, and the fish of the sea, whatever passes along the paths of the seas."*
*[ Psalm 8:5-8 ]*

*"You make known to me the path of life; in your presence there is fullness of joy; at your right hand are pleasures forevermore."*
*[ Psalm 16:11 ]*

*"Serve the LORD with gladness!*
*Come into his presence with singing!*
*Know that the LORD, he is God!*
*It is he who made us, and we are his;*
*we are his people, and the sheep of his pasture.*
*Enter his gates with thanksgiving,*
*and his courts with praise!*
*Give thanks to him; bless his name!"*
*[ Psalm 100:2-4 ]*

*"I will give thanks to Thee, for I am fearfully and wonderfully made. Wonderful are Thy works, and my soul knows it very well."*
*[ Psalm 139:14 ]*

*"Yet no one can fathom what God has done from the beginning to the end."*
*[ Ecclesiastes 3:11c ]*

*“Fear God and keep His commandments; for this is the whole duty of mankind.”*
*[ Ecclesiastes 12:13d ]*

*“I am the LORD, that is My name; I will not give My glory to another, Nor My praise to graven images. ‘Behold, the former things have come to pass, Now I declare new things; Before they spring forth I proclaim them to you.’”*
*[ Isaiah 42:8-9 ]*

*“What sorrow awaits those who argue with their Creator. Does a clay pot argue with its maker? Does the clay dispute with the one who shapes it, saying, ‘Stop, you’re doing it wrong!’ Does the pot exclaim, ‘How clumsy can you be?’”*
*[ Isaiah 45:9 ]*

*“I am God, and there is none like Me, declaring the end from the beginning and from ancient times things not yet done, saying, ‘My counsel shall stand, and I will accomplish all my purpose.’”*
*[ Isaiah 46:9b-10 ]*

*“And do not fear those who kill the body but cannot kill the soul. Rather fear Him who can destroy both soul and body in Hell.”*
*[ Matthew 10:28 ]*

*“And the one on whom seed was sown among the thorns, this is the man who hears the word, and the worry of the world and the deceitfulness of wealth choke the Word, and it becomes unfruitful.”*
*[ Matthew 13:22 ]*

*“For the Son of Man is going to come with His angels in the glory of His Father, and then He will repay each one according to what he has done.”*
*[ Matthew 16:27 ]*

*“Don’t let your hearts be troubled. Trust in God, and trust also in me. There is more than enough room in my Father’s home. If this were not so, would I have told you that I am going to prepare a place for you? When everything is ready, I will come and get you, so that you will always be with me where I am. And you know the way to where I am going.”*
*[ John 14:1-4 ]*

*“I glorified you on earth, having accomplished the work that you gave me to do.”*
*[ John 17:4 ]*

*“For since the creation of the world His invisible attributes, His eternal power and divine nature, have been clearly seen, being understood through what has been made so that they are without excuse.”*
*[ Romans 1:20 ]*

*“But who are you, O man, to answer back to God? Will what is molded say to its molder, “Why have you made me like this?” Has the potter no right over the clay, to make out of the same lump one vessel for honorable use and another for dishonorable use?”*
*[ Romans 9:20-21 ]*

*“For since in the wisdom of God the world through its wisdom did not know Him, God was pleased through the foolishness of what was preached to save those who believe.”*
*[ 1 Corinthians 1:21 ]*

*“All things are subjected to Him, then the Son Himself will also be subjected to Him who put all things in subjection under Him, that God may be all in all.”*
*[ 1 Corinthians 15:28 ]*

*“Behold! I tell you a mystery. We shall not all sleep, but we shall all be changed, in a moment, in the twinkling of an eye, at the last trumpet. For the trumpet will sound, and the dead will be raised imperishable, and we shall be changed. For this perishable body must put on the imperishable, and this mortal body must put on immortality. When the perishable puts on the imperishable, and the mortal puts on immortality, then shall come to pass the saying that is written: ‘Death is swallowed up in victory.’”*
*[ 1 Corinthians 15:51-54 ]*

*“We are destroying speculations and every lofty thing raised up against the knowledge of God, and we are taking every thought captive to the obedience of Christ.”*
*[ 2 Corinthians 10:5 ]*

*“According to His purpose, which He set forth in Christ as a plan for the fullness of time, to unite all things in Him, things in Heaven and things on Earth.”*
*[ Ephesians 1:9 ]*

*“[We] are God’s masterpiece. He has created us anew in Christ Jesus, so we can do the good things he planned for us long ago.”*
*[ Ephesians 2:10 ]*

*“The new man, which in the likeness of God has been created in righteousness and holiness of the truth.”*
*[ Ephesians 4:24 ]*

*“The new man who is being renewed to a true knowledge according to the image of the One who created him.”*
*[ Colossians 3:10 ]*

*“All Scripture is God-breathed and is profitable for teaching, rebuking, correcting and training in righteousness.”*
*[ 2 Timothy 3:16 ]*

*“Then I saw a new heaven and a new earth, for the first heaven and the first earth had passed away, and the sea was no more. And I saw the holy city, new Jerusalem, coming down out of heaven from God, prepared as a bride adorned for her husband. And I heard a loud voice from the throne saying, ‘Behold, the dwelling place of God is with man. He will dwell with them, and they will be his people, and God himself will be with them as their God. He will wipe away every tear from their eyes, and death shall be no more, neither shall there be mourning, nor crying, nor pain anymore, for the former things have passed away.’”*
*[ Revelation 21:1-4 ]*

*“This glorious city, with its streets of gold and pearly gates, is situated on a new, glorious earth. The tree of life will be there (Revelation 22:2). This city represents the final state of redeemed mankind, forever in fellowship with God: “God’s dwelling place is now among the people, and He will dwell with them. They will be his people, and God himself will be with them and be their God… His servants will serve Him. They will see His face.”*
*[ Revelation 21:3; 22:3-4 ]*

**********
**********

If you have a ‘neat’ story or some thoughts about an issue or current event that you would like me to try to respond to, I would be glad to give it a try...so, send them to me at: mbesh@comcast.net

Disclaimer: All the above jokes and inspirations are obtained from various sources and copyright are used when known. Other than our name and headers, we do not own the copyright to any of the materials sent to this list. We just want to spread the ministry of God’s love and cheerfulness throughout the world.

Mark

