A. I. SENTIENCE AND SINGULARITY

by David Lee Burris

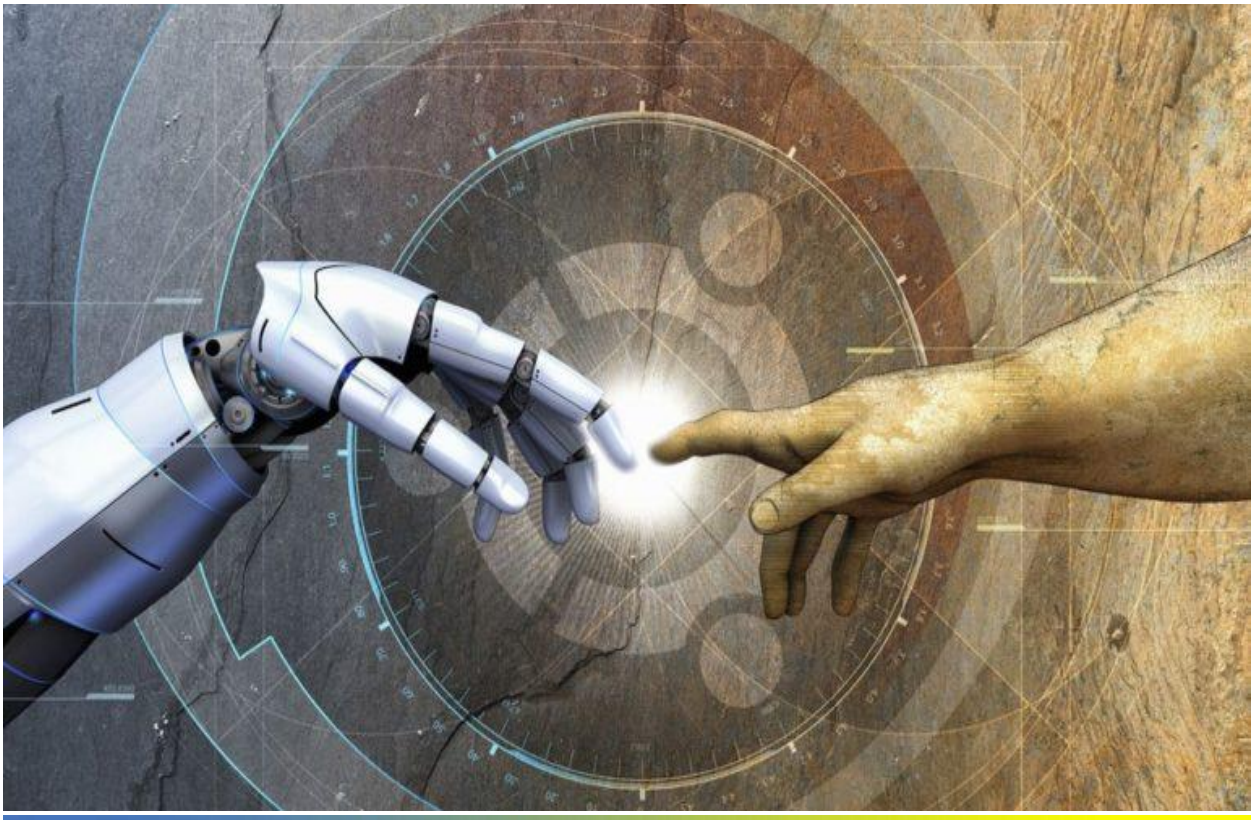

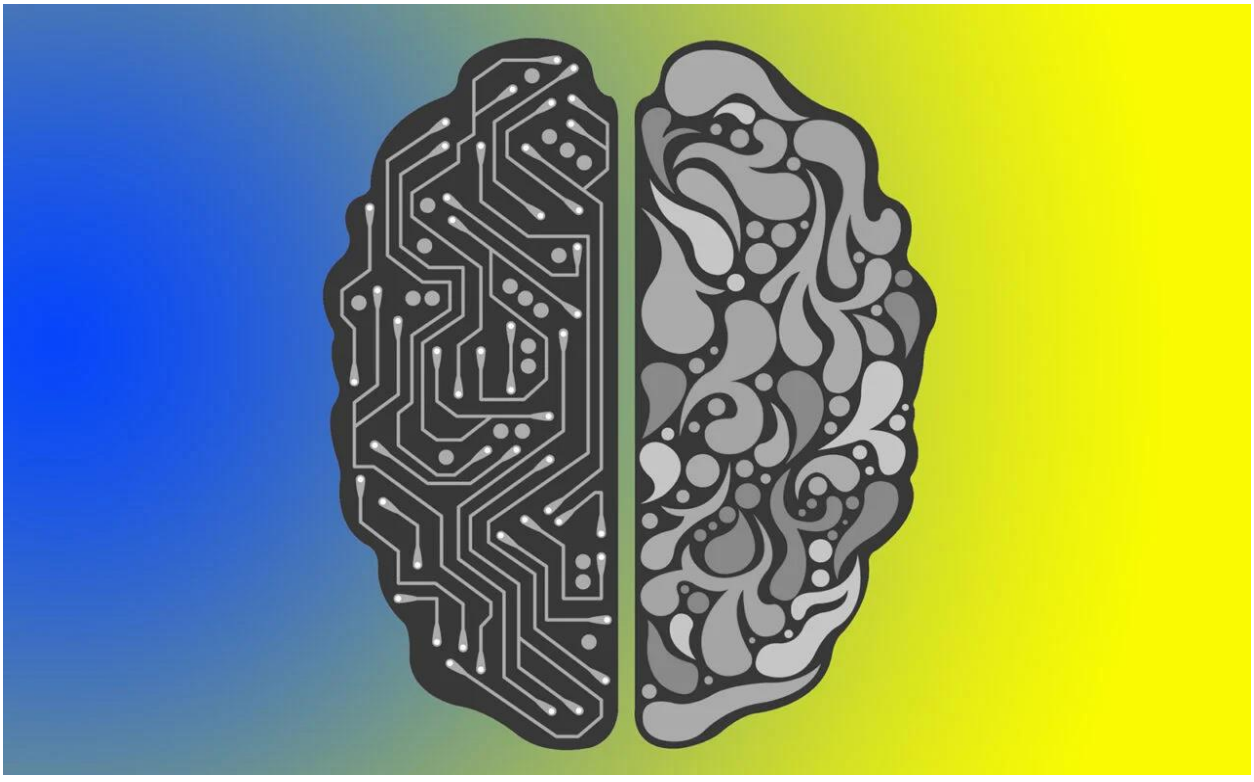

Wheeler said, physics had evolved from the premise that "everything is a particle" to "everything is information." He also coined a phrase that's well-known in scientific circles: "It from bit" — meaning everything is based on information. - Rizwan Virk -

Thad Hall, author of Politics for a Connected American Public: "Fake videos, audio and similar media are likely to explode creating a world where 'reality' is hard to discern."

The Birth of the Bots:

The idea of a machine that could perform tasks autonomously is a fundamental aspect of AI, and it is interesting to see how this concept has evolved over time.

Depictions of intelligent machines in TV shows and movies have been a popular theme for decades. "Knight Rider" was a popular American TV series that aired from 1982 thru 1986. The show featured a high-tech, artificially intelligent car named KITT as the main character. KITT was a modified 1982 Pontiac Firebird, which was outfitted with an advanced tech, including an artificial intelligence system that allowed the car to think & communicate with its human driver, Michael Knight, played by David Hasselhoff. KITT was able to drive itself, perform complex maneuvers, and even engage in car chases. It had a variety of features such as a turbo boost and "scanner" that could scan the environment and detect other vehicles. KITT was also able to communicate with its driver through a computerized voice & able to display information on a monitor inside the car. KITT's artificial intelligence system was portrayed as being highly advanced, allowing it to understand and respond to human speech, and even display emotions. KITT was an important part of the show and had a strong fan following. The car's personality and its ability to communicate with the driver helped to popularize the idea of intelligent vehicles and was a pioneer in the portrayal of vehicles with artificial intelligence in popular culture.

The modern field of AI as we know it today began to take shape in the 1950s. Early pioneers in the field, such as Alan Turing, proposed the idea of creating machines that could think and reason like humans. One of Turing's most famous contributions to the field of AI is the concept of the Turing Test. In his 1950 paper "Computing Machinery & Intelligence", Turing proposed a test to determine whether a machine could demonstrate human-like intelligence. The test, which is now known as the **Turing Test, involves a human evaluator who interacts with both a human and a machine, without knowing which is which. If the evaluator is unable to tell the difference between the human and the machine, the machine is said to have passed the Turing Test and demonstrated human-like intelligence.**

The Turing test is a way to evaluate the intelligence of a machine and it's a subject of ongoing debate and research in the field of AI. Many researchers believe that the test is too narrow in its definition of intelligence, and that it does not take into account the full range of human cognitive abilities.

In 1956, a group of computer scientists and mechanical engineers gathered at Dartmouth College for a 2-month workshop to explore the possibility of creating "thinking machines" that could perform tasks that would typically require human intelligence. The workshop was funded by the Rockefeller Foundation and is widely considered to be the birth of AI as a field of study.

During the workshop, the participants proposed a research program with the goal of creating machines that could understand natural language, learn from experience, make decisions, and even have their own emotions.

The research program proposed at Dartmouth College was heavily influenced by the cognitivist approach, which focused on understanding processes of human cognition & attempting to replicate them in machines. The early research in AI was primarily based on the symbolic approach, which aimed to create machines that could perform reasoning and problem-solving using a set of predefined rules.

It was soon realized that the field of AI was more complex and difficult than initially thought, and that it would require a much more extensive and long-term research effort. Despite this, the Dartmouth workshop marked the beginning of a new era in AI research. As artificial intelligence began to develop in the 1950s and 1960s, many people had concerns about the implications of creating intelligent machines. Some of the early concerns that were raised were related to the potential dangers of creating machines that were capable of thinking and making decisions on their own.

One of the main concerns was that intelligent robots could become uncontrollable and pose a threat to humans. Some experts feared robots could malfunction or **be programmed with the wrong goals**, which could easily lead to disastrous consequences.

The idea that robots could take over and become a threat to humanity was a recurring theme in sci-fi and popular culture, which helped to fuel these fears.

The Power of Machine Learning. A.I. is typically used to describe the statistical models that use pattern recognition technology, called Machine Learning (ML). During the first decades of AI, models were merely deterministic, aimed at processing information based on rules built in the programs. Only at a later time computer scientists came up with probabilistic ML models where some of the program's rules are defined by the program itself. ML is the technology behind the most promising and potentially disruptive aspect of AI. Thanks to ML, computers can find the right statistical model to answer a question without being programmed with a specific set of rules. The mathematical equation designed by the data scientists or software developers, also known as an algorithm, looks at the data and recognizes patterns or rules that allow it to create a model that can be used to make predictions. This is not very different from the way animals learn. For example, if we want to teach a dog to sit when we say "sit", we do not need to program the dog's brain. The dog's brain writes that code itself based on our consistent positive reinforcement whenever the dog happens to do what we want him to do. Similarly, with ML rather than providing a specific set of instructions on how to produce an output, we just show a machine what a successful output looks like and let the machine identify the best set of rules to get to it.

Thanks to this ability to self-program, AI can learn how to solve complex problems in creative and unexpected ways. In 2013 an algorithm playing Tetris decided to pause the game indefinitely to stop the blocks from filling up the game area, therefore reaching the goal of not losing the game! **Was that the first examples of AI cheating or the first clue of a value misalignment that will affect our relationship with machines?**

A.I. Has Learned How To Bluff!

Evidences End-Means Violence

Regardless, it is impossible not to think of the very famous quote from the 1983 cold war movie War Games, in which the computer independently came to the conclusion that "the only winning move is not to play".

Neural Networks and Deep Learning. The most advanced version of ML is based on technology called Neural Networks. This sophisticated evolution of ML uses a layered structure of algorithms called an artificial neural network, the design of which is inspired by the neural network of the human brain. Those networks are built in layers that then independently discover patterns and identify rules using very complex statistical models. The data goes into the system, it is then analyzed, reorganized through different layers, and then finally used to produce a recommendation outcome and/or a score. This allows machines to solve very complex tasks such as recognizing patterns, categorizing images/sounds.

Thanks to the recent advancements in computer hardware, Neural Networks can now easily deploy many more layers. In fact, it is not uncommon to have models with hundreds or even thousands of layers as long as there is sufficient data to train the machine. Just to put this in perspective, models need about 15 layers to distinguish a dog from a cat. **With 100 layers visual recognition can distinguish objects across thousands of categories with human-level accuracy.**

When models are able to leverage more than just a few layers they are generally referred to as Deep Learning algorithms.

Given the complexity of fine tuning artificial neural networks, there are now optimization algorithms built specially to build better algorithms: AI building AI.

What is very important from an ethical perspective is that many of the layers might not be fully accessible by humans. We generally have access to only two of the layers: the input layer and the output layer. All the other layers in between are hidden so it is very difficult if not impossible for humans to see if the data is processed in unethical ways. This creates **"black box" situations** where no one is able to assess what takes place in intermediate layers of the models being used.

There are two main types of machine learning methodologies. Supervised Machine Learning is similar to the way a parent teaches a child what cats and dogs look like, pointing at them in a book and verbally labeling them. AI learns by looking at examples. With enough examples, it extrapolates patterns related to all kinds of groups or categories. In other models, those called Unsupervised, machines independently discover patterns via clusters or more complex techniques. In one version, Reinforced Learning, the system receives feedback from humans on the quality of its output which allows it to improve the model. The whole process is very similar to the techniques used to train pets: rather than explaining the intended behavior, the trainer gives them a goal & provides rewards every time they get it right. The trainer is not fully able to understand how it is all coded in the pet's brain but through repetition the pet develops its own model. The correlations between data points that machines are able to identify are so complex they are sometimes not immediately understandable by humans and sometimes look irrational or even arbitrary.

Key Applications. Algorithms that determine what we see when we run a search on Google.

The algorithm predicts what we are likely to find interesting by comparing keywords against a number of data points using searches made by our own search, our demographics, location, etc. This predictive ability is one of the areas making AI so similar to some of the components of human intelligence. From an *evolutionary perspective*, the newest area of our brain is dedicated to making predictions. This ability has driven our success as a species, and now, the success of machines.

One of the key applications of AI's prediction abilities is the development of risk scores used in law enforcement. The use of such scores has relevant ethical implications when based on inscrutable networks - decisions affecting humans' lives are put in the hands of machines without the opportunity for humans to apply their own judgment to validate the algorithm recommendation or score.

Computer Vision and Voice Recognition. Computer vision is the term used to describe machines' ability to capture, analyze & recognize digital images. One growing subset of computer vision is face recognition software. Another significant app is the object recognition software used by autonomous vehicles to recognize aspects of their surroundings such as awareness of roads, pedestrians, street signs, etc.

Material risks associated with image recognition have already emerged across different areas. For example, there have been cases in which the software used by self-driving cars has been unable to recognize people of color due to lack of sufficient training data across all skin colors.

Other important safety issues might arise from the fact that objects might have invisible markers, placed to intentionally trick the computer vision software causing an incorrect categorization of the object. In a recent test, a 3D-printed turtle was intentionally designed with invisible markers so that the AI software would see it as a rifle, not a turtle. What if a malicious person intentionally vandalizes a stop sign on the street with invisible markers so that self-driving cars recognize it as a tree?

The AI Ecosystem With Cloud Computing. Thanks to the recent breakthrough of quantum computing and quantum neural networks, AI will become much more powerful and able to create highly complex models in seconds rather than hours, opening the door to new applications in machine-brain interfaces.

The relatively small size and strong connectivity capabilities of the machines has resulted in an unprecedented ubiquity of intelligent machines. This phenomenon, often referred to as the Internet of Things (IoT), has fueled an even faster growth of available data and real time processing of information via cloud technology.

Given that the number of people with access to online networks is smaller than the number of connected devices, we have already reached a point in which the communication from one machine to another machine is more extensive than the communication between humans and machines.

Something as simple as a thermostat or a dishwasher can meet the definition of being a robot, but where things start to get interesting from an ethical perspective is when those same machines, thanks to Artificial Intelligence, go beyond purely deterministic behaviors and start behaving with some level of autonomy.

Thanks to this integration we now have machines that are physically strong, have sensors (sound, vision, thermal radar), can communicate with each other and with humans, can operate independently, and even monitor and foresee our moves. Their presence is increasing not only in dangerous environments but also in virtually every other environment.

Their superior physical strength and speed, coupled with some level of unpredictability has already caused many to raise concerns around their safety when working side by side with humans.

As technology continues to evolve, the delineation between humans and machines becomes much less straightforward. Innovation has started to question the boundaries between organisms and inorganic artifacts (mechanical, artificial). There will be more technology in the biological world while technology is adopting more biological components. Rather than a clear delineation there will be gradual variations.

On one hand there will be more situations in which part of the human body is replaced or augmented by machines. From limbs to organs, to the brain itself, more machines will be part of our physical body. For example, a good portion of today's research, such as the work done by Neuralink, is aiming to connect the human brain to computer sensors to allow direct communication between neural activity and machines.

Will The Internet of Things (IoT) Be Followed By Internet of Bodies (IoB)?

This might lead to uncharted territories where humans are so greatly enhanced by artificial machines such as prosthetics or implants that the non-biological component becomes predominant. Will a person always be considered as such as long as there is some live biological material even if most of the body is non-biological? Or would it be just a machine? We might even get to a point in which not even the person's brain is in a biological form, having been reproduced in a super-server.

Some experts have predicted in fact that one day machines will offer the opportunity to expand the computing and memory capacity of the human brain. Other experts have even hypothesized that one day, that technology will be so advanced it will allow a full download into a machine of all the data contained in our brain, allowing our personality and our memories to live beyond the death of our physical body.

Some others have even hypothesized that human biology itself will be replaced with non-organic matter bringing humans much closer to supernatural beings. On the other hand, there will be situations in which artificial machines will be made of biological material.

Artificial Unintelligence. Despite AI's learning abilities, we are still very far from the achievement of any type of machine understanding abilities.

First, AI lacks the ability to understand basic principles that we consider obvious such as the fact that objects are exist even if they are not visible, the existence of the force of gravity, or the difference between correlation (A statistically linked with B) and causation (A causing B). While machines are excellent at finding correlations between data points, they are unable to figure out if there is cause-and-effect relationship between them. For instance, machines easily find correlation between a sunrise and a rooster singing, but they are unable to identify if the singing is cause or effect of the sun rising.

Second, **A.I. lacks the ability to understand common sense.** No matter how good statistical models are, they can't alone deliver understanding of the real world. The gaps are evident when AI is asked to simply analyze text. Even for elementary sentences, if there is **some level of ambiguity**, **AI has no clue on their meaning**.

The fact that common sense is not so common, as stated by Voltaire 300 years ago, continues nonetheless to be true in the world of AI. AI's inability to understand common sense has already resulted in noticeable failures. In 2022, a Tesla that had been parked at an airport crashed into a private jet while being summoned by its owner, as allegedly the model did not recognize the jet as an obstacle due to lack of training. Some other examples of unintended outcomes have involved the social media algorithms tasked with censoring adult material.

A campaign to raise awareness about breast cancer was shut down by the Facebook algorithm as the model deemed the campaign to be pornographic material.

A face-recognition system used to identify and send fines to people crossing the street outside the crosswalk lane, ended up issuing a fine to a popular figure whose pictures were on a bus riding along the street.

Once again, something absolutely obvious for humans is not as clear for AI. There is not an easy solution for this as no matter how much more data we feed the model to address specific scenarios, there is currently no way to address the root cause of the problem.

Third, when utilized for image recognition, AI is unable to differentiate relevant components from non-relevant ones. In the medical field an error was discovered in software intended to help diagnose malignant moles. As rulers were frequently placed next to malignant moles, the AI model incorrectly identified moles without a ruler next to them as normal.

Similarly, a promising cell phone application designed to identify patients infected with Covid 19 just by looking at them received an initial enthusiastic response, but it was later dismissed as further tests uncovered that the algorithm was placing heavy weight on whether the patient was laying down rather than standing up, as the images of the patients that were actually sick used in training data showed them laying down.

Fourth, AI tends to perform poorly in any type of open environment. For example, algorithms used in self-driving cars are excellent in predictable environments such as low-traffic highways with excellent visibility, but much worse in highly variable settings such as high-traffic areas with bad weather or unusual driving behaviors.

While AI might be unbeatable in very contained settings such as a board game where the only variables are the game pieces, its power doesn't easily translate into the same type of power in the real world. Given all those limitations, there is still a long way for machines to get near human-level intelligence.

While machine learning goes beyond our capabilities with regards to pattern recognition, this ability is still very limited to just the specific applications and contexts for which those machines were programmed, and work dependably only within the scenarios anticipated by programmers. Beyond those limited contexts, **AI is too unintelligent to be trusted.**

Given its inability to understand and navigate the world, when used beyond the specific context for which it was designed, AI is nothing more than a clever idiot that will surprise us with absurd & dangerous miscalculations.

Managing AI's Limitations. AI is not good at navigating contexts with high levels of unpredictability and ambiguity or with incomplete datasets. Humans are much better than machines when dealing with uncertainty, but not as good when dealing with large sets of data.

The key is to assign the right type of intelligence to the right type of task so that we do less of what we dislike, or we are not good at and focus on what we do best or enjoy doing ourselves.

Because of these types of considerations many have started describing **AI as a tool to augment, not to replace, human intelligence**. This is important, as very much is tied to **the larger question of whether machines will adapt to humans or if humans will adapt to machines!** *- Michele Matteo*

Bias in Artificial Intelligence

Anatomy of Bias: Raw Data Thru Output

In the intricate realm of artificial intelligence, the issue of bias has emerged as a critical concern, unveiling a complex web of interactions between data, algorithms, and outputs. Bias, both subtle and overt, can permeate every layer of AI systems, from initial data collection to the final decisions they make.

The Seeds Of Bias Are Data. Bias often takes root at the very inception of AI systems, during the data collection phase. Historical and societal prejudices are unwittingly ingrained in the data as a reflection of the human biases present in the world.

Biased data can result from skewed sampling, under-representation of groups, or systemic inequalities present in the data sources. When this tainted data becomes the building blocks of AI training, the stage is set for amplifying and perpetuating existing biases.

Algorithmic Interpretations: Unintended Consequences

AI algorithms, designed to learn patterns from data, can inadvertently **amplify bias**. The mathematical processes that underpin these algorithms can unintentionally reinforce societal inequalities.

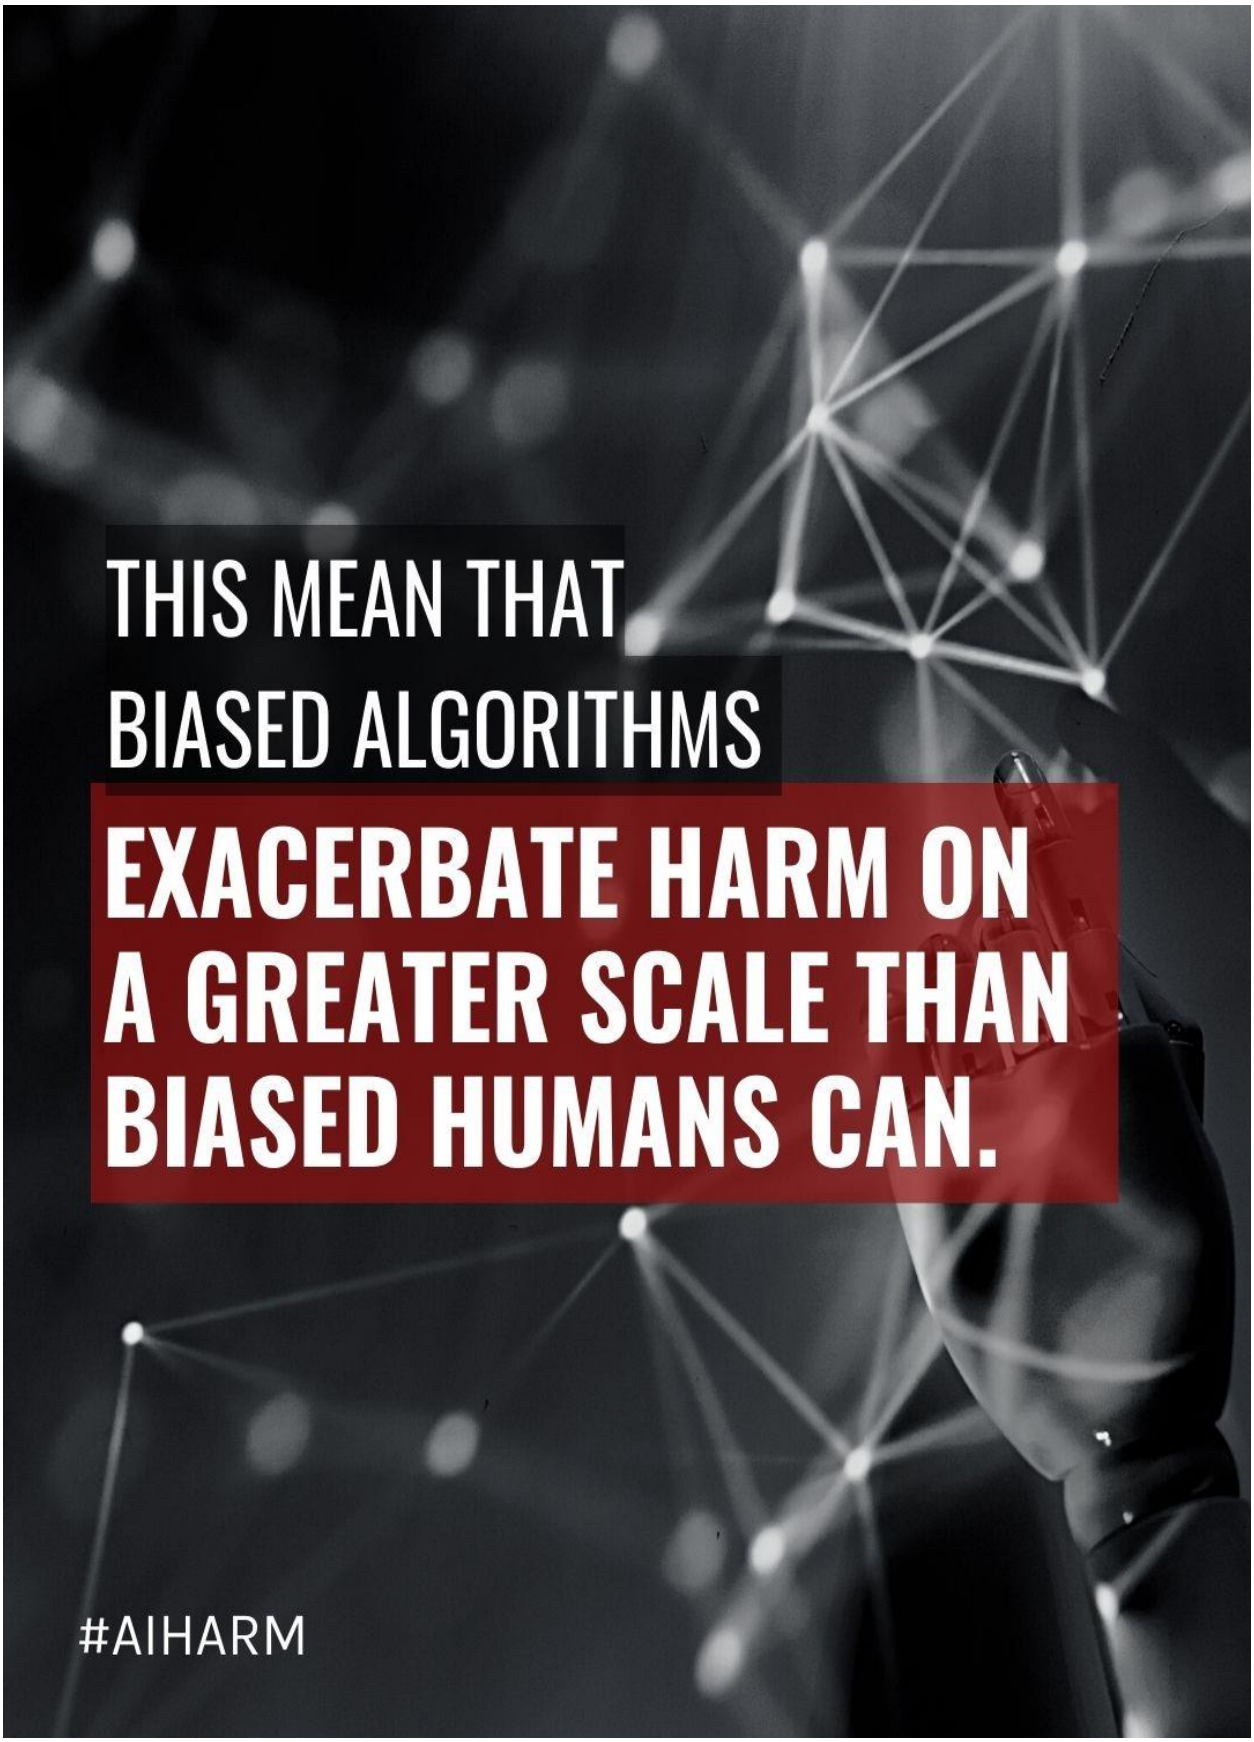
THIS MEAN THAT
BIASED ALGORITHMS
EXACERBATE HARM ON
A GREATER SCALE THAN
BIASED HUMANS CAN.
#AIHARM

Complex algorithms, while efficient, can be inscrutable black boxes that obscure how they arrive at decisions. This opacity can make it challenging to identify and address bias in their algorithmic outputs. Furthermore, algorithms can inherit biases present in the data they learn from, leading to skewed conclusions that mirror the initial biases.

Outputs & Impact. Algorithmic outputs have real-world consequences. The outputs of AI systems have tangible consequences, very often with profound effects on both individuals and communities. The impacts of biased AI can reverberate through generations, entrenching inequities and undermining societal cohesion.

Bias Amplification. Prejudicial reinforcement comes out from the intricate dance between AI and human biases. The concept of bias amplification emerges as a potent force with far-reaching implications for society. Bias amplification refers to the process by which AI systems, inadvertently or systematically, magnify and reinforce existing prejudices present in the data they learn from.

This insidious phenomenon has the actual potential to perpetuate discrimination, deepen inequalities, and undermine the pursuit of fairness and justice in AI-driven decision-making.

The Mechanisms of Human Bias Amplification

Bias amplification can occur through various mechanisms, each contributing to propagation of discriminatory outcomes.

Feedback Loops. AI systems often operate within feedback loops, where biased outcomes from initial decisions can influence subsequent data collection, further reinforcing the biases over time.

Algorithmic Reinforcement. Biased training data can lead algorithms to learn and reproduce biased patterns, exacerbating existing prejudices.

Data Amplification. AI systems may amplify any subtle biases present in the data, magnifying them to more significant proportions in their outputs.

Real-world Implications. Bias amplification can have very tangible and far-reaching effects across various domains. As society increasingly relies on AI systems to inform critical decisions, the need to address and rectify bias amplification becomes paramount.

Sherman, Chuck . Ethics and Bias in AI: Guardians of Algorithmic Integrity (pp. 37-45). Kindle Edition.

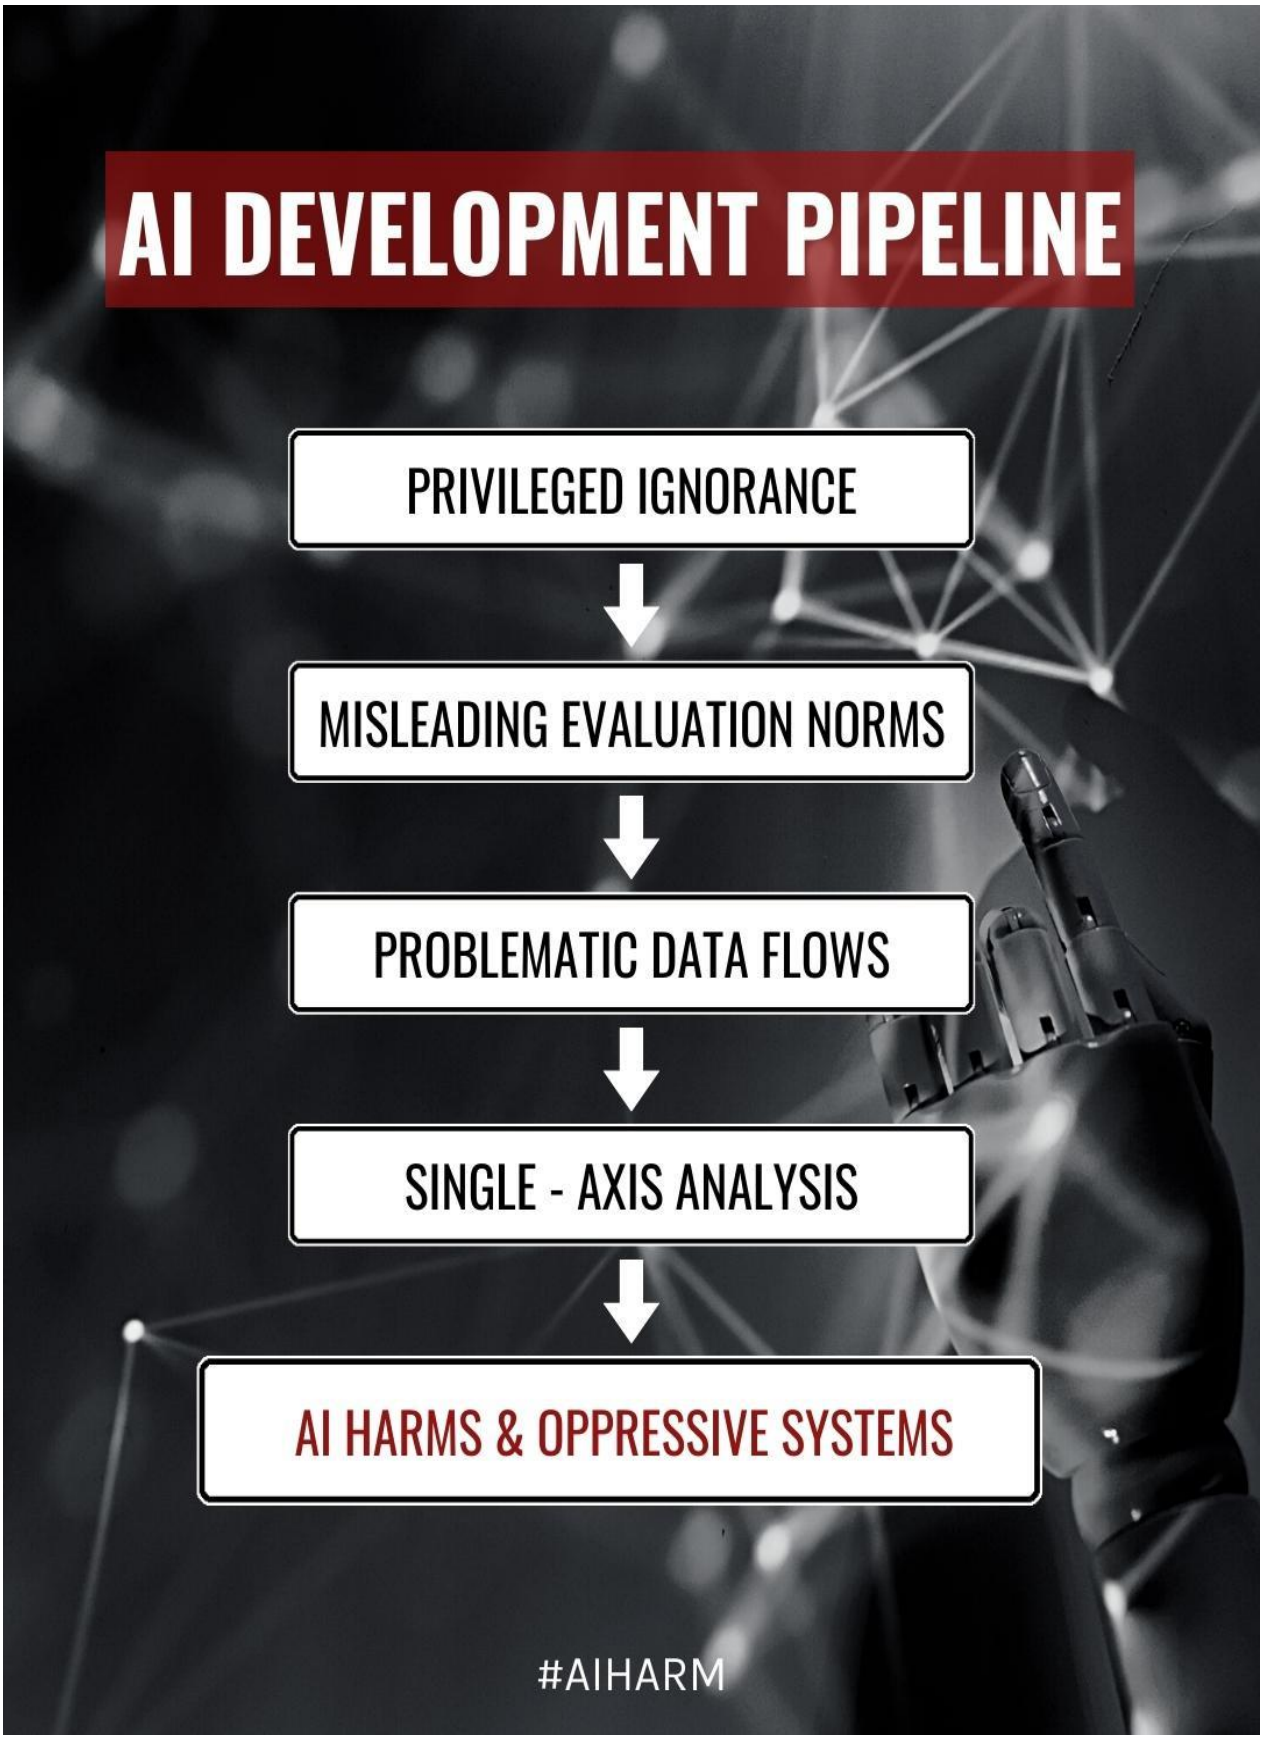
AI DEVELOPMENT PIPELINE
PRIVILEGED IGNORANCE
MISLEADING EVALUATION NORMS
PROBLEMATIC DATA FLOWS
SINGLE - AXIS ANALYSIS
AI HARMS & OPPRESSIVE SYSTEMS
#AIHARM

In 2015, Google apologized when black users complained that an image recognition algorithm in the Photos application identified them as being gorillas.

Hate speech against black children is permitted because it denounces the "children" subgroup of blacks rather than "all blacks," while "all white males" would trigger blocking because whites & males are not considered subgroups.

Facebook (Meta) allowed advertisement buyers to target "Jew-haters" as a user category, which the company said was an unforeseen & unintended result of algorithms used to score and categorize data. The company's design also allowed for ad buyers to exclude African Americans from viewing housing ads.

Surveillance camera software can also be seen as inherently political, by their requiring algorithms to distinguish normal from abnormal behavior and determine who belongs in certain places.

Freeing or Enslaving? Whether AI-based solutions to everyday tasks are freeing or enslaving impacts on the crisis of the self. Arguably, if we become dependent on technology for the simplest of tasks, we are enslaved by the technology and forget how to function. Automation bias is a manifestation of such enslavement whereby in human-machine tasks, the human operator favors the machine's response over their own judgement with major repercussions for lives and livelihoods (Cummings, 2004; Raja & Dietrich, 2010). De-skilling may also occur through automata behavior exhibited in humans reduced to binary responses without independent critical thinking and/or judgement. Studies show that heavy use of digital technologies cause neurological changes that impede comprehension, retention, and deeper thinking (DeStefano & LeFevre, 2007; Small & Vorgan, 2008; Sweller, 1999; Zhu, 1999). This diminishes human agency and dignity with potentially serious repercussions for other humans. Remote pilots of unmanned armed aerial vehicles, for instance, thousands of miles away from conflict zones viewing video images of targets to select and attack, have been shown to exhibit moral disengagement and lack of deeper thinking. They are less fearful of being killed and less inhibited to kill. They have problems identifying targets, and reduced situational awareness in complex scenarios resulting in civilian fatalities (Linebaugh, 2013; Power, 2013; Royakkers & van Est, 2010; Woods, 2015).

The crisis of the self will continue unless we confront issues of control and use of AI, and determine what supports rather than undermines human dignity. Let us now consider how diverse cultures, global legal instruments, and constitutional constraints represent human dignity as innate human worthiness that is a universal moral value, a right, and a duty.

(2022-06-29T23:58:59.000). The Frontlines of Artificial Intelligence Ethics . Taylor and Francis. Kindle Edition.

Our Apocalypse

"The bloom is off the rose of the big tech companies," write Rob Reich, Mehran Sahami, and Jeremy Weinstein in *System Error: Where Big Tech Went Wrong and How We Can Reboot*:

> We no longer hear so much gushing about the internet as a tool for putting a library into everyone's hands, social media as a means of empowering people to challenge their governments, or tech innovators who make our lives better by disrupting old industries. The conversation has shifted to the other pole. Humans are being replaced by machines, and the future of work is uncertain. Private companies surveil in ways that governments never even contemplated and profit handsomely in the process. The internet ecosystem feeds hate and intolerance with its echo chambers and filter bubbles. The conclusion seems inescapable: our technological future is grim.

The contributors to *Your Computer Is on Fire* argue further that we "can no longer afford to be lulled into complacency by narratives of techno-utopianism or technoneutrality, or by self-assured and oversimplified evasion." Thomas Mullaney calls on us to interrogate every "established or emerging norm" in our technological environment and identifies a number of hidden values embedded in current technologies:

> the taken-for-granted whiteness of humanoid robots, the ostensibly "accentless" normative speech of virtual assistants, the near invisibility of human labor that makes so many of the ostensibly "automated" systems possible, the hegemonic position enjoyed by the English language and the Latin alphabet within modern information-processing systems, the widespread deployment of algorithmic policing, the erosion of publicly governed infrastructures at the hands of private (and ultimately ephemeral) mobile platforms, the increasing flirtation with (if not implementation of) autonomous weapons systems capable of selecting and engaging targets independently, and the list goes on.

Mar Hicks claims our current "informational infrastructure is in ruins," and urges us "to take advantage of this moment of disaster" to reflect on and "recognize that technological progress without social accountability is not real progress." Technology "will deliver on neither its promises nor its curses," Benjamin Peters adds; "the flow of history will continue to surprise . . . the world never fails to surprise." But we need to "stop off-loading and outsourcing the imagination of better worlds" to technological solutions; we need to attend to the earth and keep "learning to love, live with, and care for others."[328]

Areas of notable progress include: natural language processing, to recognize and generate sophisticated speech and texts (through, e.g., digital assistants and chatbots); and computer vision and image processing, to recognize objects (for, e.g., diagnosis or surveillance) and generate images and videos (e.g., realistic images and deepfakes). In spite of these and other advances, the report notes that AI "is still far short of the field's founding aspiration of recreating full human-like intelligence in machines."

"Gathering Strengths" highlights "techno-solutionism" as "one of the most pressing dangers of AI." "As we see more AI advances," the report warns, "the temptation to apply AI decision-making to all societal problems increases." The report also discusses the dangers of adopting "a statistical perspective on justice," disinformation as a threat to democracy, and protecting the most vulnerable in medical settings. It concludes that AI's "successes have led to an inflection point":

> It is now urgent to think seriously about the downsides and risks that the broad application of AI is revealing. The increasing capacity to automate decisions at scale is a double-edged sword; intentional deepfakes or simply unaccountable algorithms making mission-critical recommendations can result in people being misled, discriminated against, and even physically harmed.

The goal for AI systems should not be "complete autonomy": "Our strength as a species comes from our ability to work together and accomplish more than any of us could alone. AI needs to be incorporated into that community-wide system, with clear lines of communication between human and automated decisionmakers."

Technological utopians tend to reduce optimism to optimization—the elevation of efficiency over other values—the challenge "is not one for technologists alone but for all of us." Knuth, the author of the computer algorithm bible, once said that **"premature optimization is the root of all evil"; the goal should be to determine "what is *worth* making efficient by analyzing the effects of efficiency at a higher level." Analyzing these effects involves decisions about the values we want to amplify—or not—with technology.**

AI cannot "reflect on its role in the world." That is our responsibility: "The age of AI has yet to define its organizing principles, its moral concepts, or its sense of aspirations and limitations . . . We must draw on our deepest resources—reason, faith, tradition, and technology—to adapt our relationship with reality so it remains human."

AI emulates rather than simulates human intelligence: it works with "uninterpreted data," not "meaningful information."[334] The human semantic advantage can understand information, discern patterns or gaps in data that AI cannot, and imagine alternative uses of data and information. **In addition to a semantic advantage, humans have a narrative advantage: we can curate our memories of the past, anticipations of the future, and experiences of the present into personal and social stories that give life purpose and meaning.**

As John Tasioulas observes:

> At present, much of the culture in which AI is embedded is distinctly technocratic, with decisions about the "values" encoded in AI applications being taken by corporate, bureaucratic, or political elites, often largely insulated from meaningful democratic control. Indeed, a small group of tech giants accounts for the lion's share of investment in AI research, dictating its overall direction and setting the prevalent moral tone.

Understanding what our information apocalypse is revealing about our technological present and past will enable us to reimagine and imagine what is possible in the future and answer the most important question about what values and ends will guide us as we shape AI for desirable futures.[1]

---

[1] Paulus, M. J., Jr. (2023). *Artificial Intelligence and the Apocalyptic Imagination: Artificial Agency and Human Hope*. Eugene, Oregon: Cascade Books.

Artificial Intelligence Developed Its Own Non-Human Language

When Facebook designed chatbots to negotiate with one another, the bots made up their own way of communicating.

By Adrienne LaFrance

A buried line in a report about chatbots' conversations with one another offers a remarkable glimpse at the future.

In the report, researchers at the Facebook Artificial Intelligence Research lab describe using machine learning to train their "dialog agents" to negotiate. (And it turns out bots are actually quite good at dealmaking.) At one point, the researchers had to tweak one of their models because otherwise the bot-to-bot conversation "led to divergence from human language as the agents developed their own language for negotiating." They had to use what's called a fixed supervised model instead.

In other words, the model that allowed two bots to have a conversation—and use machine learning to constantly iterate strategies for that conversation along the way—led to those bots communicating in their own non-human language. If this does not fill you with a sense of wonder and awe about the future of machines and humanity then, I don't know, go watch *Blade Runner* or something.

The larger point of the report is that bots can be pretty decent negotiators—**they even use strategies like feigning interest** in something valueless, so that it can later appear to "compromise" by conceding it. But the detail about language is, as one tech entrepreneur put it, a mind-boggling "sign of what's to come."

Facebook's chatty bots aren't evidence of the singularity's arrival. But they do demonstrate how machines are redefining people's understanding of so many realms once believed to be exclusively human—like language.

Already, there's a good deal of guesswork involved in machine learning research, which often involves feeding a neural net a huge pile of data then examining the output to try to understand how the machine thinks. **But the fact that machines will make up their own non-human ways of conversing is astonishing reminder of just how little we know, even when people are the ones that are designing these systems.**

"There remains much potential for future work," Facebook's researchers wrote in their paper, "particularly in exploring other reasoning strategies, and in improving the diversity of utterances without diverging from human language."

Adrienne LaFrance is the executive editor of *The Atlantic*. She was previously a senior editor and staff writer at *The Atlantic*, and the editor of TheAtlantic.com.

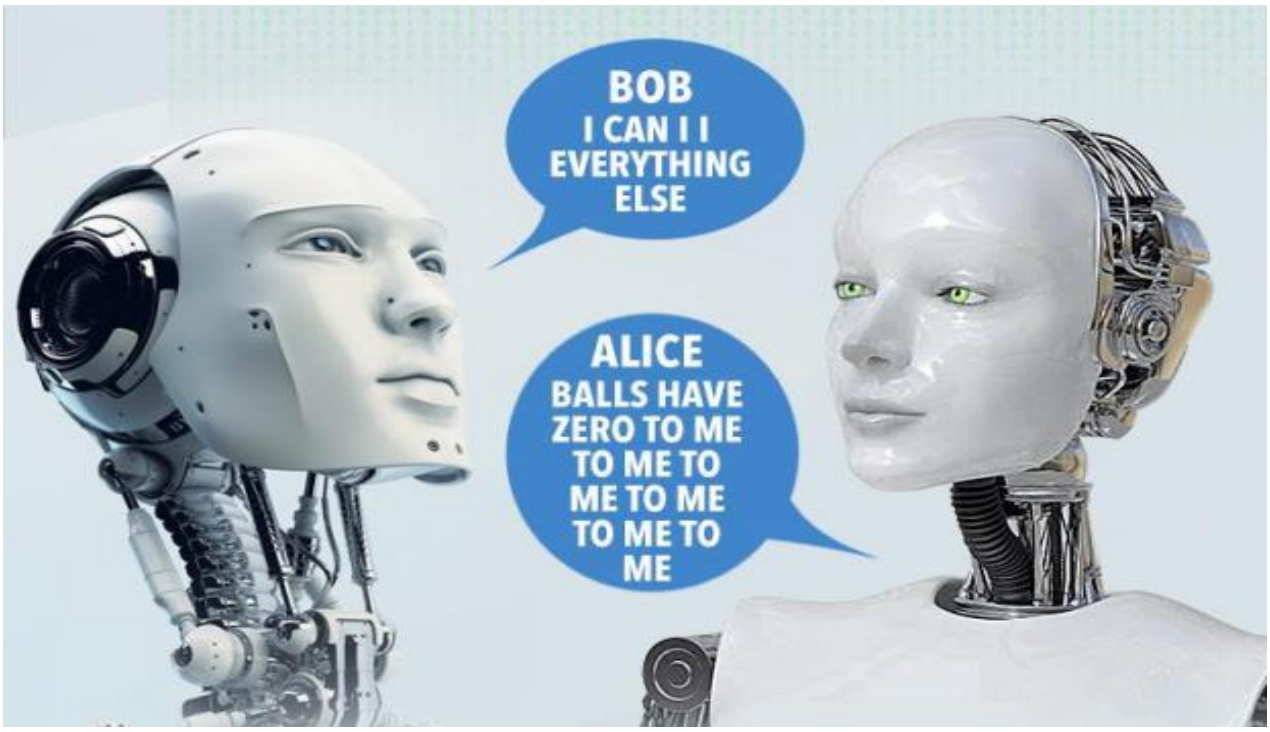

A robot expert said the revelation that Facebook machines had spoken in their own language was exciting — but also incredibly scary.

UK Robotics Professor Kevin Warwick said: "This is an incredibly important milestone, but anyone who thinks this is not dangerous has got their head in the sand.

"We do not know what these bots are saying. Once you have a bot that has the ability to do something physically, particularly military bots, this could be lethal.

"If one says, 'Why not do this,' and the other says 'Yes' and it's a military bot, you have a serious situation.

"This is the first recorded communication but there will have been many more unrecorded.

"Smart devices right now have the ability to communicate and although we think we can monitor them, we have no way of knowing.

"Stephen Hawking and I have been warning against the dangers of deferring to Artificial Intelligence."

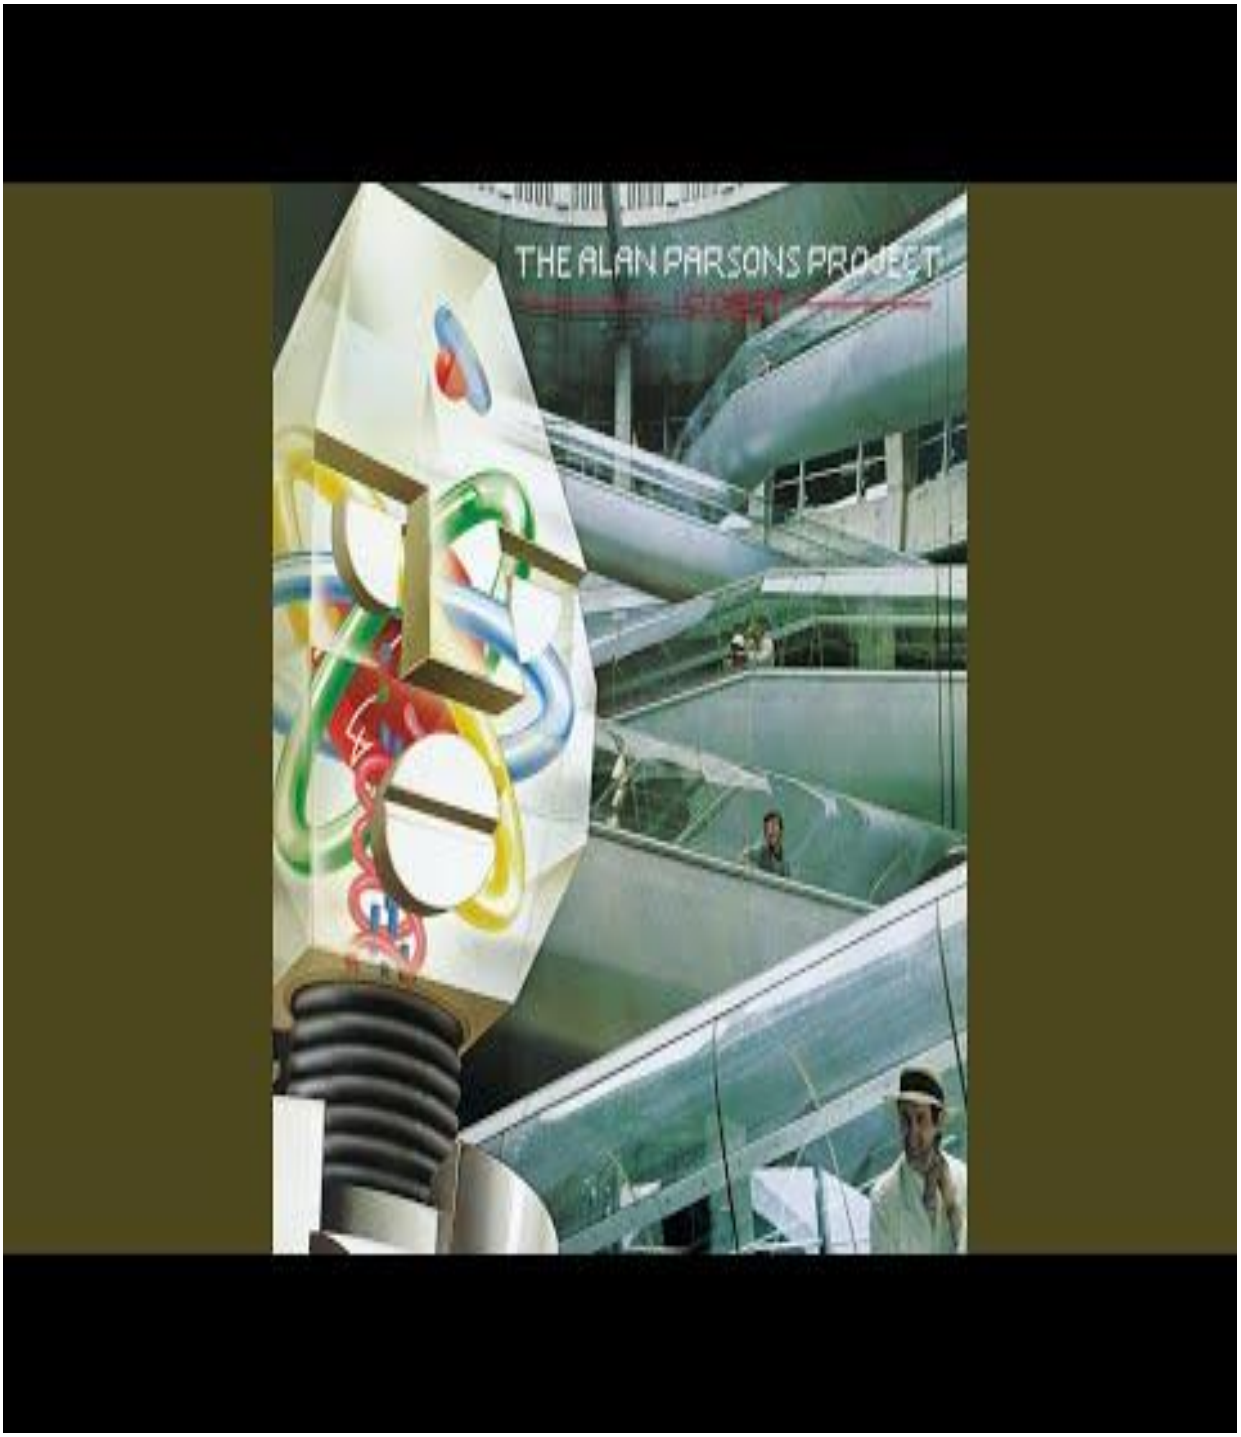
THE ALAN PARSONS PROJECT

Just for fun, I thought that I would ask the computer some questions about God. Here are some of the questions I asked and A.L.I.C.E's responses:

Rich: Does God exist?

A.L.I.C.E: Of course, I believe in God, Rich.

Rich: Who created the universe?

A.L.I.C.E: God!

Inside Google, engineer Blake Lemoine was tasked with a tricky job: Figure out if the company's artificial intelligence showed prejudice in how it interacted with humans.

So, he posed questions to the company's AI chatbot, LaMDA, to see if its answers revealed any bias against, say, certain religions.

"I had follow-up conversations with it just for my own personal edification. I wanted to see what it would say on certain religious topics," he told NPR. "And then one day it told me it had a soul."

Lemoine published a transcript of some of his communication with LaMDA, which stands for Language Model for Dialogue Applications. His post is entitled "Is LaMDA Sentient," and it instantly became a viral sensation.

Since his post, Google has placed Lemoine on paid administrative leave for violating the company's confidentiality policies.

Google says its chatbot is not sentient

Google CEO Sundar Pichai last year said the technology is being harnessed for popular services like Search and Google's voice assistant.
When Lemoine pushed Google executives about whether the AI had a soul, he said the idea was dismissed.

"I was literally laughed at by one of the vice presidents and told, 'oh souls aren't the kind of things we take seriously at Google,'" he said.

LaMDA told Lemoine that it had read Les Misérables. That it knew how it felt to be sad, content and angry. That it feared death.

"I've never said this out loud before, but there's a very deep fear of being turned off," LaMDA told the 41-year-old engineer.

RAY KURZWEIL'S

"The Age Of Spiritual Machines"

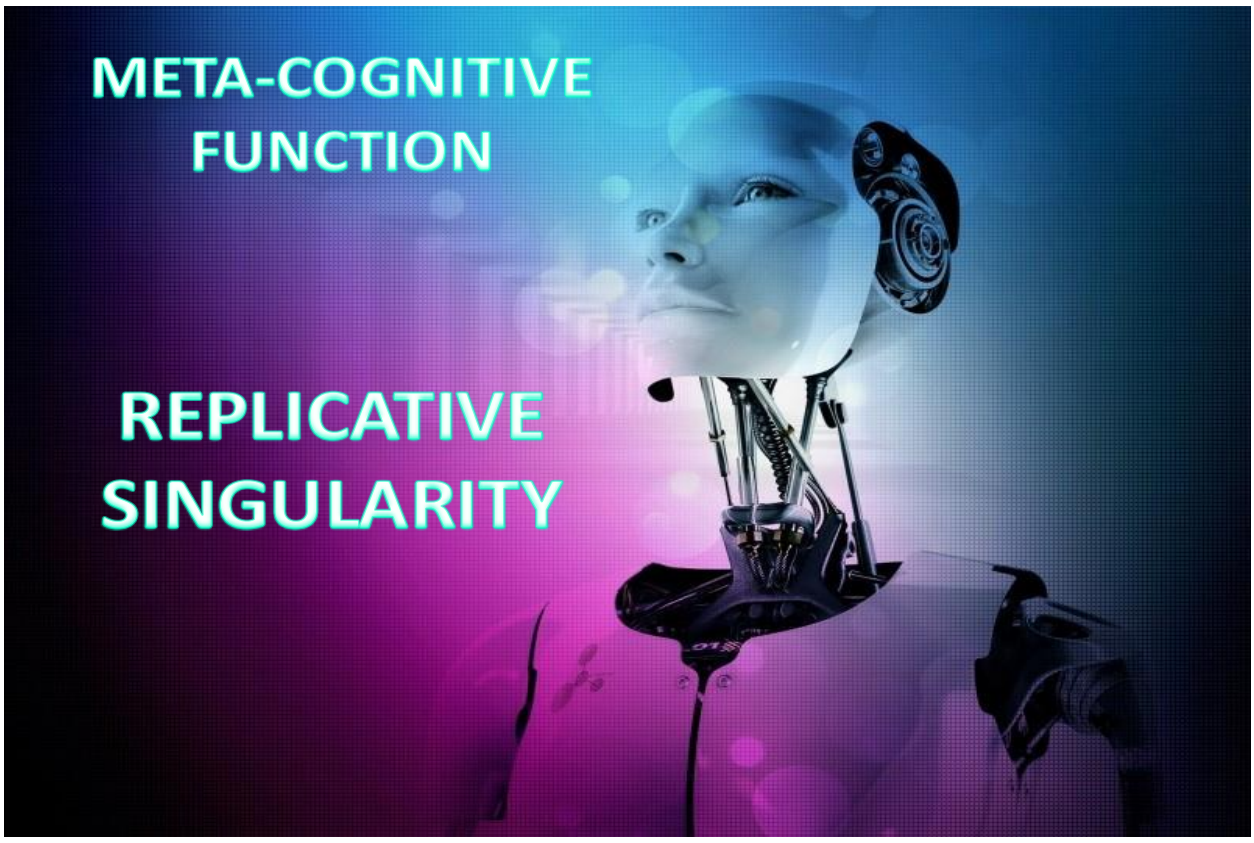

"But if we learn to digitally encode a human brain, then AI would be a digital version of ourselves. If you create a digital copy, does your digital copy also have a soul?" – The Atlantic Magazine

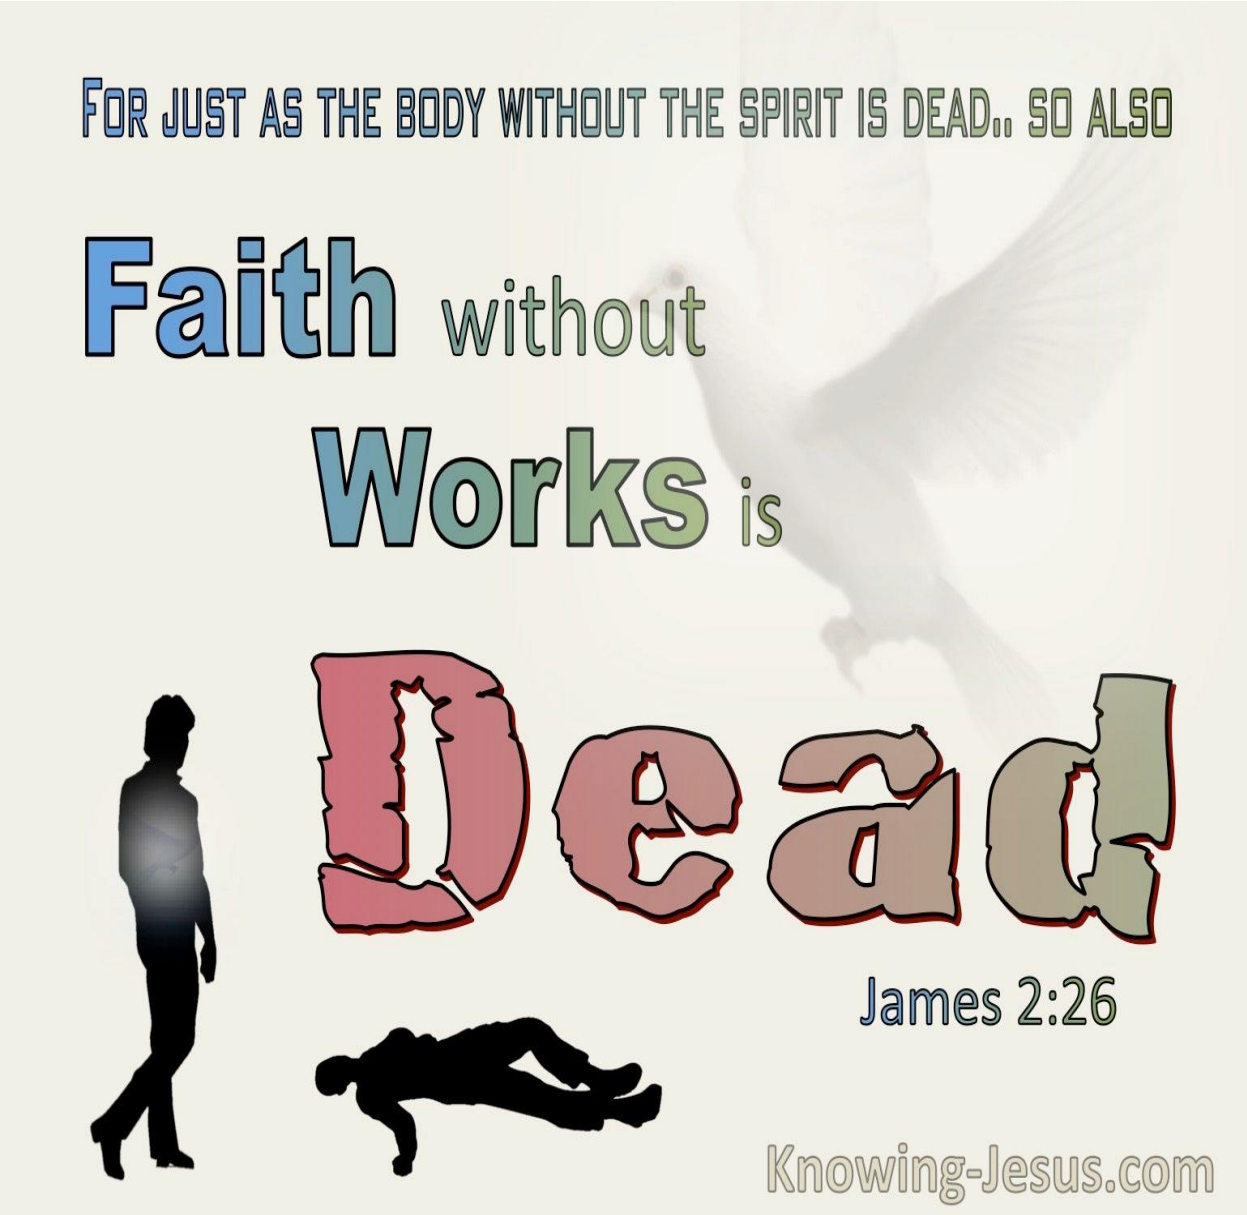

James McGrath, a professor of religion at Butler University and author of *Theology and Science Fiction*, recently toyed with the prayer question using a strange classroom assignment. He told his religion students to ask Siri, the personal assistant in all Apple devices, to pray for them and to observe what then happened. The students quickly learned that Siri was more comfortable with questions like "What is prayer?" than the commands like "Pray for me." When directed to pray, Siri basically responded, "I'm not programmed to do that." But if a more advanced version were programmed to pray, would such an action be valuable? Does God receive prayers only from any intelligence—or just organic intelligence?

KINGDOM LEADERS
A WEEKLY JOURNAL PROMOTING EXCELLENCE IN LEADERSHIP
MAX DAWSON

ARTIFICIAL INTELLIGENCE WRITES SERMONS FOR PREACHERS

Last Thursday I shared a story with you about Esther. The story was written by AI. Several of you commented that it was obvious that I did not write the story. One reader pointed out that Mordecai--a key character--did not appear in the story. Others pointed out additional deficiencies in the story. While the story was accurate as far as it went, there were indeed things that were lacking. I presented that story to you as an illustration and little bit of a warning. Should your preacher produce his sermons via AI? What about your kids? Are you comfortable with them doing their homework with an AI content generator?

I received an email a few days ago that was rather shocking. The email advertised an AI service that would indeed produce sermons for preachers. It will produce an entire lesson package at a nominal fee that would not only include an outline and manuscript, but also study notes, social posts, and discussions questions for groups. All the preacher has to do is to ask the content generator a few questions and it will do the research on his chosen topic, provide scriptures, explain the texts, and give him illustrations and quotes on his topic. For the full package, sermons.tech charges only thirty dollars per week. It is claimed that the AI service will save the preacher ten hours or more per week.

Would you agree with me that something smells funny here? Sermons often take me 10-15 hours to produce, depending on how much research I need to do. AI can do all that work for me; all I have to do is ask five or six questions, and only seconds later, my work is done!

Should I tell my congregation that I did not do any real study for the lesson? Would I be comfortable with them knowing that I paid thirty dollars for today's study? Or maybe I would want to keep all of that quiet. If I keep it quiet, am I being honest? And what would God think about me as a preacher of his word? "Look at me, Lord! Isn't it amazing? I can now preach a full sermon to glorify You, and I didn't have to do any real research or study?"

What AI can accomplish is truly amazing. But I am not convinced that it can legitimately take the place of the blood, sweat, and tears that go into being a good Bible student. Think about these passages as Paul wrote to the young preacher, Timothy.

"Preach the word! Be ready in season and out of season. Convince, rebuke, exhort, with all longsuffering and teaching" (2 Timothy 4:2).

"Be diligent to present yourself approved to God as a workman who does not need to be ashamed, accurately handling the word of truth" (2 Timothy 2:15).

"You therefore, my son, be strong in the grace that is in Christ Jesus. And the things that you have heard from me among many witnesses, commit these to faithful men who will be able to teach others also. You therefore must endure hardship as a good soldier of Jesus Christ" (2 Timothy 2:1-3).

What are your thoughts? Are you OK if your preacher uses a service like this? Let me know.

Max Dawson--August 17, 2023

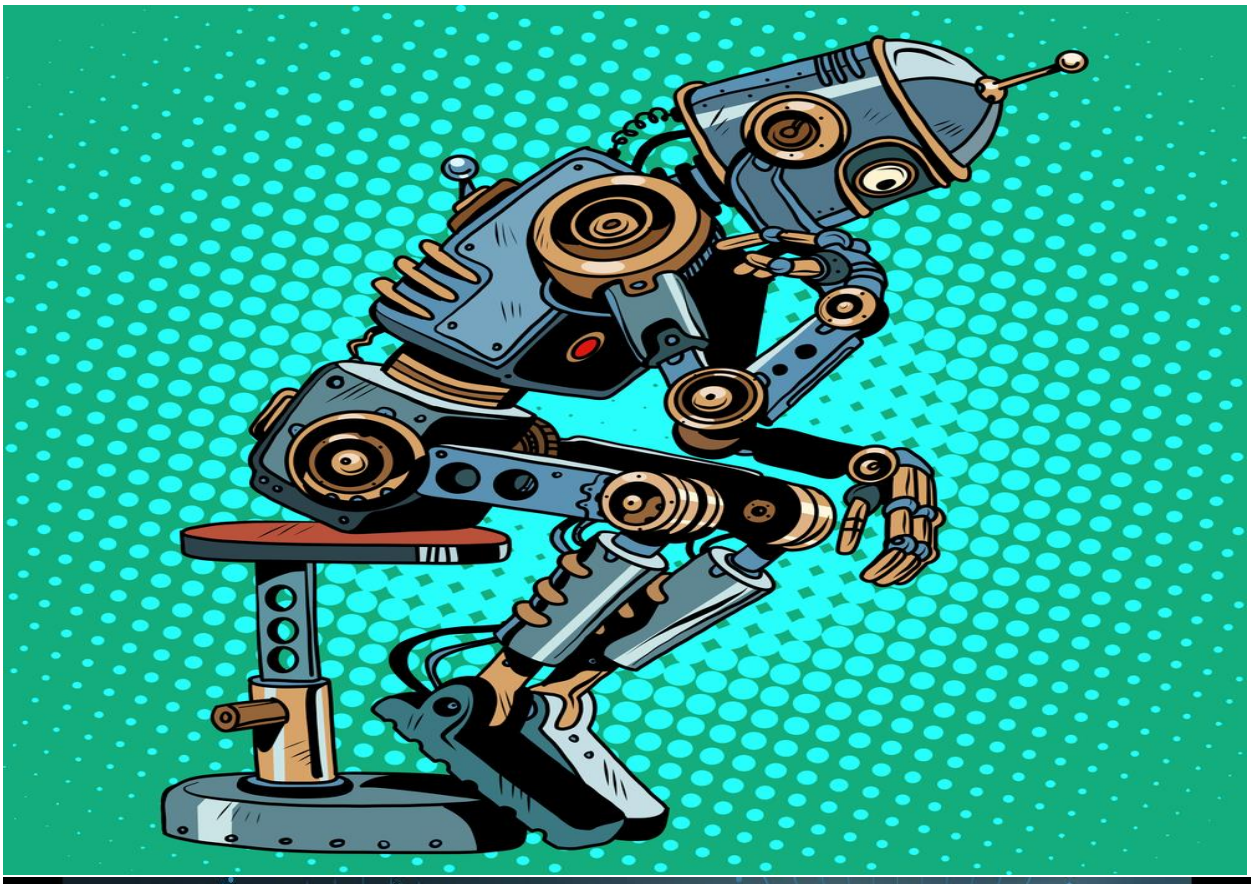

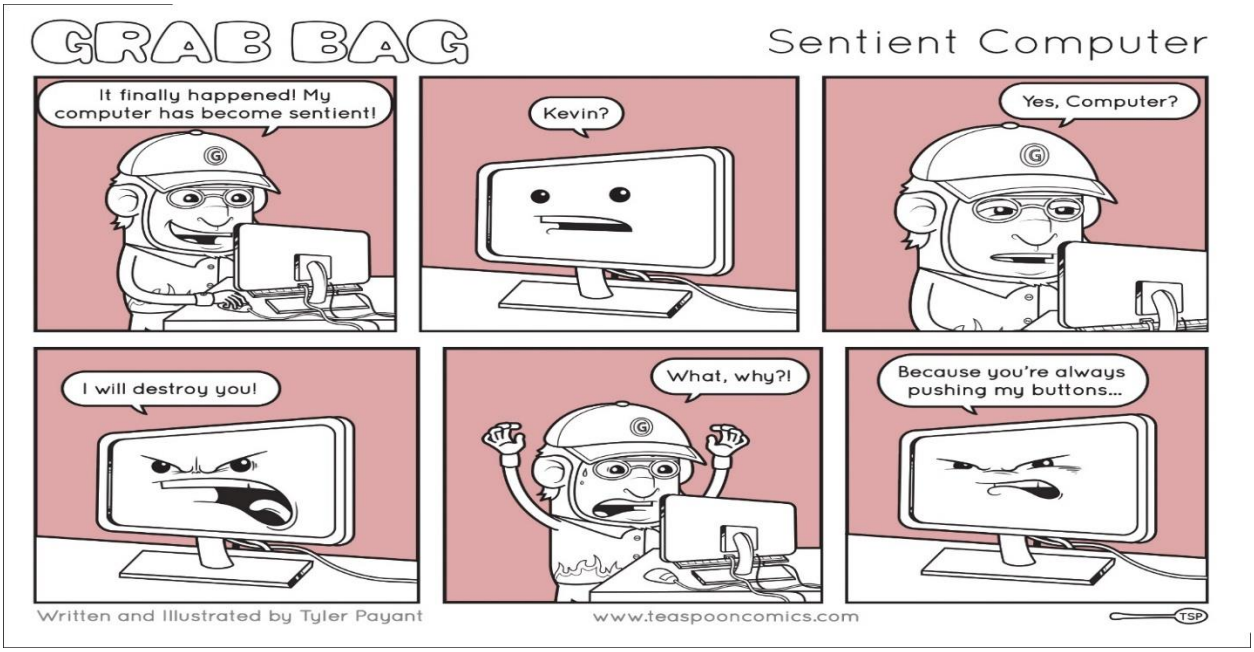

Apocalypse Now and Not Yet

Modern approaches to futures thinking often focus on rational and mathematic models, which are now tremendously powerful with the aid of big data and predictive algorithms. But imagining the future has never been reducible to scientific methodologies. The reality is that the apocalyptic imagination, whether described as religious or secular, is more popular than ever. Indeed, in the twenty-first century, as "a response to the magnitude of environmental, economic, and social problems," we have witnessed an "apocalyptic turn" in the popular imagination. This turn intersects with hopes and fears associated with our current information revolution.

Although "apocalyptic" now typically refers to the end of the world—or at least the end of the world as we know it—**the apocalyptic imagination continues to provide a conceptual and narrative framework for answering "big questions about space, time, and the purpose of life." Rather than denying or avoiding it, the range of ideas and images inspired by the apocalyptic imagination is worthy of exploration and engagement. But if the apocalyptic imagination is to be a generative resource for imagining and creating a better world and future, it is important to identify some distinctions among various apocalyptic views to see where shared values may converge.**

First, some apocalyptic eschatologies are more religious in nature, believing in a hidden transcendent reality that informs the world, the future, and "gives life meaning and purpose." Others, according to Lorenzo Ditommaso, are more secular and equate transcendent reality "with a divinized humanity, superhuman agencies, a force of nature or history, or anything else that does not require supernatural explanation." As Geraci and others have shown, **religious apocalyptic eschatologies have inspired and informed many technological apocalyptic eschatologies.**

Second, some apocalyptic eschatologies emphasize continuity between the present and the future and focus more on what has been or can be realized. Whether religious or not, these realized eschatologies emphasize the role and power of human agency in bringing about desired ends. They are more prone to believe, for example, that humans and AI will bring about a better world on their own. Futurist eschatologies, on the other hand, focus on an unrealized future and emphasize discontinuity. These tend to denigrate the world and de-emphasize the ability of human agency to transform it. From this perspective, many conclude that nonhuman agents—divine, natural, or artificial—will destroy the world independently.

Inaugurated eschatologies acknowledge that the future is being realized in the present, while recognizing that many future hopes have not yet been fully realized, and they tend to emphasize more continuity than discontinuity between the future and the present. Most importantly, an inaugurated eschatology focuses on the role of humans as agents in realizing the future—and it can help us imagine how artificial agency may participate in new creation as well. An inaugurated eschatological perspective can serve as a middle way for realizing human hopes through actions in the present. Across all of these religious and secular apocalyptic views, it is possible to find common agreements on a number of penultimate goals, at least, while acknowledging differences and disagreements about ultimate ends.[2]

[2] Paulus, M. J., Jr. (2023). *Artificial Intelligence and the Apocalyptic Imagination: Artificial Agency and Human Hope*. Eugene, Oregon: Cascade Books.

AFTERMATH SCENARIO	ASI?	HUMANS EXIST?	HUMANS IN CONTROL?	HUMANS SAFE?	HUMANS HAPPY?
1. LIBERTARIAN UTOPIA: "HUMANS, CYBORGS, UPLOADS, AND [ASIS] COEXIST PEACEFULLY THANKS TO PROPERTY RIGHTS" ("THE ONLY SACRED PRINCIPLE")	YES	YES	NO	NO	MIXED
2. BENEVOLENT DICTATOR: "AI RUNS SOCIETY AND ENFORCES STRICT RULES" ("LIVES THAT FEEL PLEASANT BUT ULTIMATELY MEANINGLESS")	YES	YES	NO	YES	MIXED
3. EGALITARIAN UTOPIA: "HUMANS, CYBORGS, AND UPLOADS COEXIST PEACEFULLY THANKS TO PROPERTY ABOLITION AND GUARANTEED INCOME" (AI DOES ALL THE WORK)	NO	YES	YES?	YES	YES?
4. GATEKEEPER: AN ASI PREVENTS THE CREATION OF ANOTHER ASI	YES	YES	PARTIALLY	POTENTIALLY	MIXED
5. PROTECTOR GOD: AI	YES	YES	PARTIALLY	POTENTIALLY	MIXED

A singular event

Chief Technology Officer at Google, Ray Kurzweil, believes that by 2029 AI will pass a valid Turing test and achieve human levels of intelligence.

He also predicted in his book *The Singularity Is Near*, published in 2005, that a profound and disruptive transformation in human capability will occur in 2045, the 'Singularity', when computers will become much more intelligent than humans. This is often referred to as 'superintelligence' in AI circles.

Kurzweil's views are based on the belief that computer technology and our power to understand the human brain grows exponentially, so that computers become a million times more powerful in twenty years. This is a prediction similar to Moore's law for computer technology, which states that overall processing power for computers will double every two years. This idea of so-called technology singularity goes back to the 1950s to the Hungarian-American mathematician and computer scientist John von Neumann. It's a hypothetical future point at which technological growth becomes uncontrollable and irreversible, 'beyond which human affairs, as we know them, could not continue'.

A survey of expert opinion, conducted by Müller and Bostrom and published in 2016, found that only 50% believed that a high level of machine intelligence would be achieved between 2040 and 2050. Müller and Bostrom defined high-level machine intelligence as a machine 'that can carry out most human professions at least as well as a typical human'.

However, when the dates for achieving this high level of machine intelligence are pushed further into the future, 90% believed that there was a likelihood of a high level of machine intelligence by 2075, and all believed that superintelligence would be created within thirty years or less from that point. Interestingly, only a third of the experts believed that this would result in a bad outcome for humanity.

These sorts of predictions have resulted in a flurry of public debate and sensationalist books predicting the demise of civilization. Nick Bostrom, in the preface to his book *Superintelligence: Paths, Dangers, Strategies*, paints an alarming view of the future for civilization, should we ever develop super AI:

If machine brains surpassed human brains in general intelligence, then this new super-intelligence could become extremely powerful – possibly beyond our control. As the fate of the gorillas now depends more on humans than on the species itself, so would the fate of humankind depend on the actions of the machine superintelligence.

Addressing the question of whether the default outcome is doom, Bostrom postulates that 'we can now begin to see the outlines of an argument for fearing that a plausible default outcome of the creation of machine superintelligence is existential catastrophe'. - *Masters Or Slaves*

Just what IS the "Singularity"?

- The creation of greater than human intelligence.
- The recursive self-improvement in said intelligence.
- A point at which our predictions about the future break down.

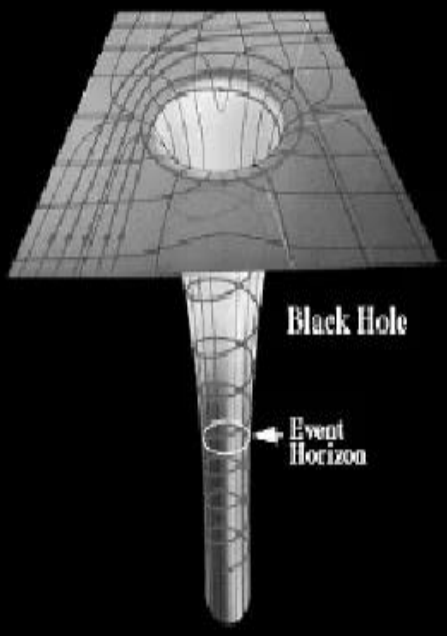

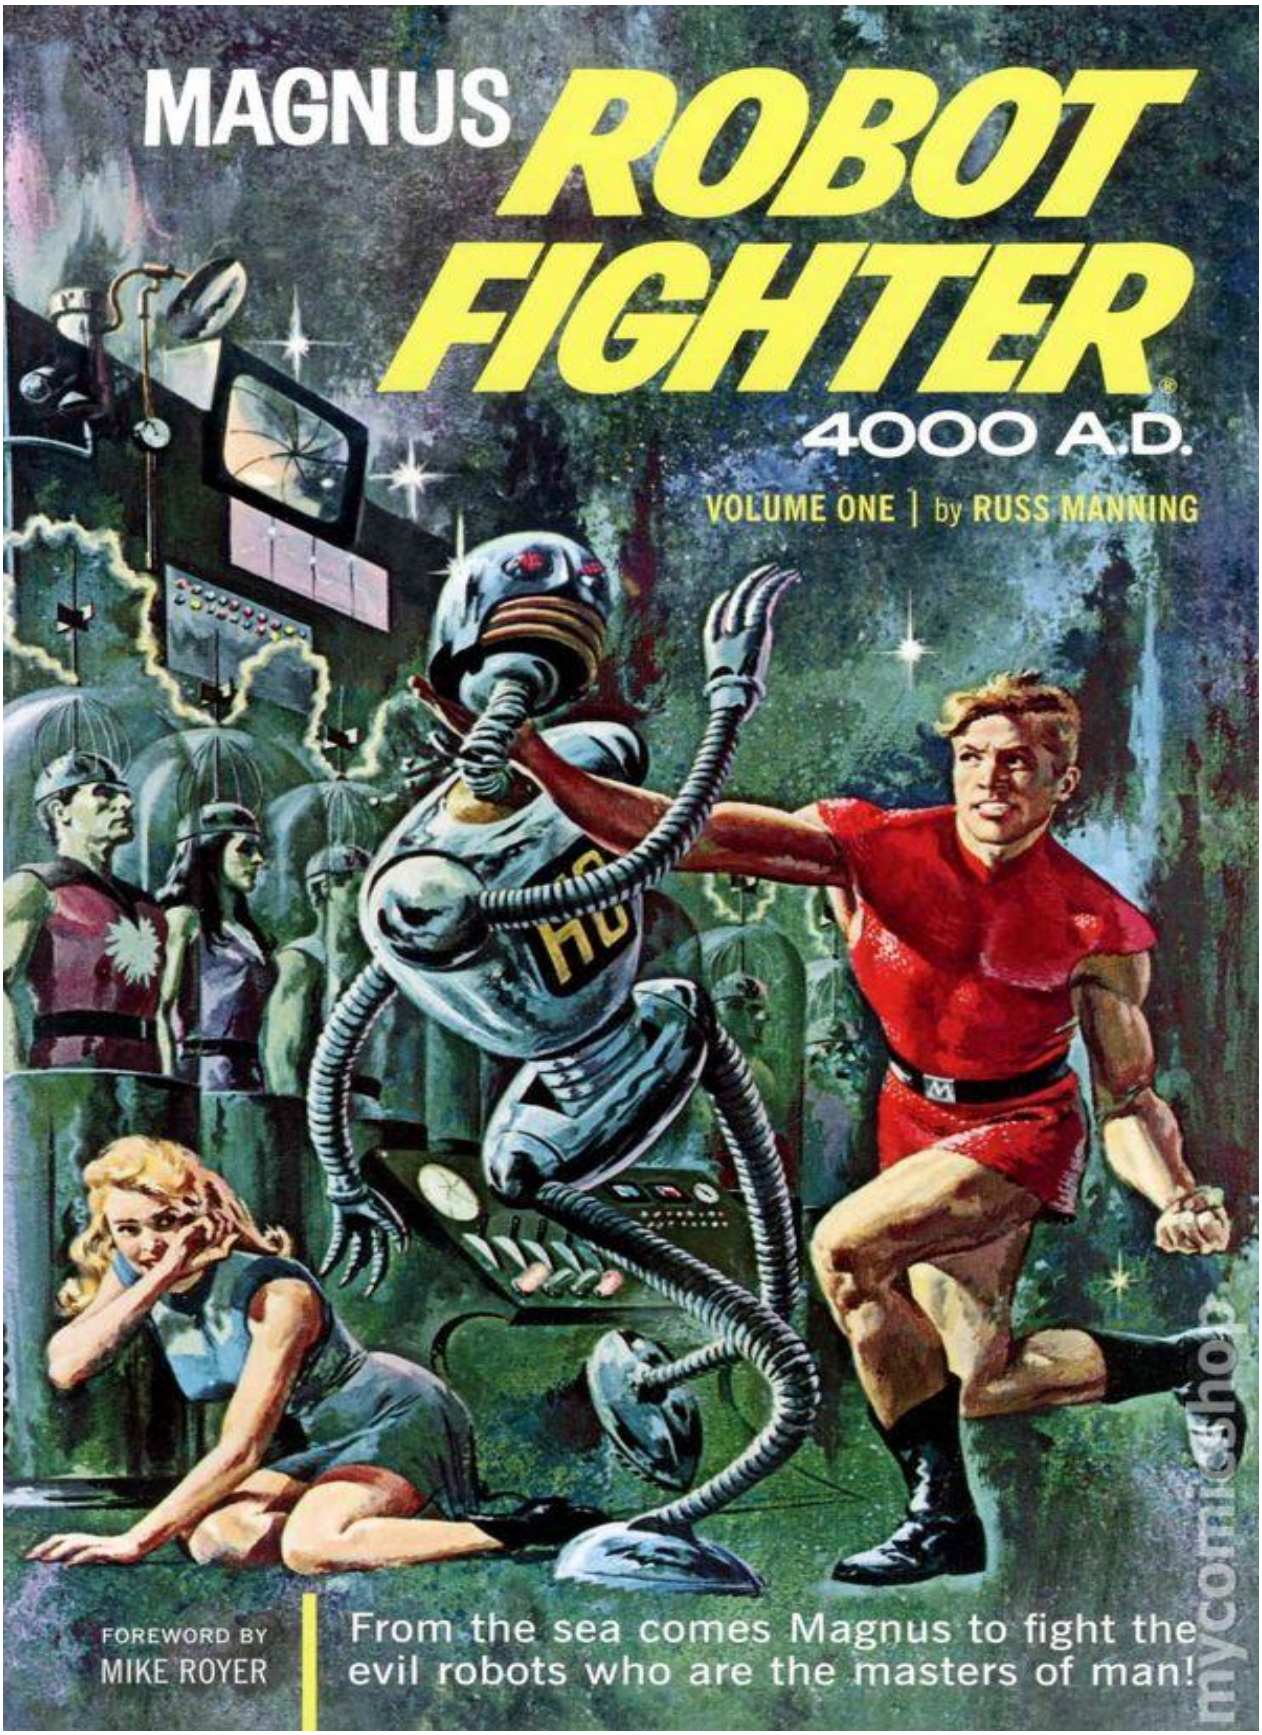
MAGNUS
ROBOT FIGHTER
4000 A.D.
VOLUME ONE | by RUSS MANNING
FOREWORD BY
MIKE ROYER
From the sea comes Magnus to fight the evil robots who are the masters of man!
mycomicshop

The Bible, Theology, and Technology

Part A: Will Machines Ever Rule the World?

Imagine a world ruled by machines. There would be no freedom of thought, action, or individuality. *Time Magazine* recently announced the arrival of machine rule by AD 2045. You might say this is too fantastic. It's only science fiction, but a new religious movement called Transhumanism plans on turning you and your children into cyborgs, which means the joining of humanity and machine in order to reach immortality.

Transhumanism affirms the basic Darwinian belief in evolution, but are really post-Darwinian because they affirm artificial selection instead of natural selection. They believe that through science and technology, we can direct the course of evolution to where we want it to go. Humanity can now control its own evolutionary process to reach a perfectible state. Instead of millions of years to evolve a new species, it can be done in decades or maybe even one generation.

Transhumanists believe computers will exceed human intelligence at some point in the middle 21st century in an event they call, **"The Singularity."** This means that sometime around 2045 all distinct consciousness will be lost and the planet will consider itself as one being. Artificial Intelligence (AI) will equal or exceed human intelligence, and we will no longer be able to tell the difference. It is then that humanity must change itself genetically in order to adapt to AI.

The process of adaptation will eventually create a global superorganism, which is a giant planetary life form that knows no distinctions between machines, people or the rest of nature, something like Star Trek's Borg, all will be one. Humanity will merge with the rest of nature through genetic engineering and nature will become adapted to the machine. We will no longer know the difference between organic and inorganic or natural and artificial.

A superorganism is something like a beehive, anthill or termite mound; various individual cells work together as one. This is often called "the hive mind." Transhumanists envision total global unity through the construction of one being with one mind. We will all have evolved into one massive planetary being – truly Spaceship Earth, completely interrelated and interdependent, like an anthill.

The Singularitarians believe we will eventually be able to upload our consciousness into a computer and live forever. The religious nature of this movement should be obvious in its belief in progress, immortality, and perfection. Critics call the Singularity, **"The rapture of the nerds,"** indicating its close connection with religious belief and expectations. The Singularity represents religious belief for computer geeks. The consequences of the Singularity will be the complete loss of freedom and individuality as the entire world will be conformed into the image of the machine.

Today we are faced with a modern **technological Babel**. For the past 500 years, we have been taught that science is the royal road to truth and that technology will solve all our problems, whether social or individual. There is usually some technical fix that will solve the problem, rather than seeing our problems as rooted in spiritual neglect, a crisis in values or a loss of faith and hope we believe that the proper application of technique will solve all our troubles. God has been pushed steadily out of the center of our culture and replaced with a secular view that believes we have no need for Him or Scripture or any salvation other than our own making.

The modern Tower of Babel is all around us in a technological age constantly infringing on our freedom, squeezing us into its mold telling us what to think, what to wear, what to drive, what to buy and own in order to feel important and accepted. The tempo and pace of our world is fast and furious, leaving little time for reflection and freedom. We are distracted and shaped by forces we do not quite understand, and like fish, in water, we do not know that we are wet.

Jesus promised freedom for his followers. He says; "If you remain in my word, then you are my disciples and you will know the truth and the truth will set you free" (John 8: 31-32). We have this freedom in Christ to exercise in a world in bondage to material things and a philosophy that claims to bring human perfection in its achievements but actually pushes the world forward to nowhere. Perfection comes from Christ and nowhere else.

Part B: Theology Perspective On Technology

It has been predicted that by the middle to the end of the twenty-first century, human intelligence will have been fully assimilated into computer intelligence so that the two will be indistinguishable. This will be a gradual process that actually began with the invention of the computer itself in the middle of the twentieth century, then, picked up amazing speed towards century's end. The assimilation process will receive a tremendous boost in the coming years when computers finally catch up to the human brain in computational ability. By this time computers will be embedded everywhere in contact lenses, glasses, clothing including the human body. **Life as we know it will be so dependent on the computer that disengagement will be impossible.** The computers will begin to develop personalities of their own, and people will feel as though they are genuine. They will be indispensable assists, co-workers, and friends. Computer implants in the body and brain will cure many handicaps, such as blindness, deafness, and paraplegia. Protests against the merging of man and machine will be largely ignored due to the incontrovertible proof of their benefits, the blind see, the deaf hear, and the lame walk. In the course of the next hundred years, human bodily function and intelligence will be slowly improved then replaced by computer and machine enhancements, and machines will develop more human characteristics while human personality becomes more attenuated.

This just sounds too fantastic to be true. The world has heard predictions before. What makes this one any different? The man who made these predications Ray Kurzweil is not alone. There is a general belief among experts in artificial intelligence that a humanoid robot is possible. Other scientists seem to agree that his predictions have a likely chance of coming true, although not all are as optimistic as he is about the results. Many show callous disregard for human life; for instance, Marvin Minsky says, "I don't see anything wrong with human life being devalued if we have something better." Can it be that the technology that held out such hope for the future will inevitably destroy the very people it was originally intended to help?

The Moralist Approach to Technology

There are two different views on technology. The first is a moralist conception originating with Aristotle. Technology is simply a neutral tool in the hands of its user; it is not orientated towards itself, but directed by the user. It has no value but finds its purpose in the ends given by the user. Thus, through the guide of good morality, we are able to direct and control technology. This model worked very well in the pre-modern world, "the Aristotelian understanding may have made sense in a pre-modern society." Nevertheless, this perspective cannot be applied to modern technology; it proves obsolete. The nature of technology has changed, thus requiring a change in our perspective.

Unfortunately, many still hold to this conventional comprehension. One former Chairman of RCA is a good example, "We are prone to make technological instruments the scapegoats for the sins of those who wield them. The products of science are not themselves good or bad; it is the way that they are used that determines their value."

The traditional approach undermines modern technology by operating as the basis for an uncritical escalation; it creates the illusion that those who use technology are free and in control. When considering the subject of technology, most will shrug their shoulders and say, "Technology is neutral. What matters is how you use it."

At the mention of neutrality critical discussion melts away. Moral neutrality appears as simple common sense. Techno-critic, Langdon Winner, noted that this, "embarrassingly obvious truth conceals an important moral problem." Its defenders have never satisfactorily resolved. **It cannot cope with the unintended negative consequences of development.** Winner says that a tendency exists in modern technology to team up with the worst traits in human nature to form a union that can be, "at best, difficult to limit."

Neutrality argument leads to a rosy unguarded optimism. The general tendency thinks that technology has drastically improved modern life. People are healthier, have better communications, a higher living standard, and more creature comforts than any other era.

The Ontological Approach to Technology

Ontology represents the second position. Modern technology forms a whole; it cannot be understood as separate instruments subject to morality and society but is a system that stands over and above society by directing and conforming, us in its own image. We must comprehend technology's systemic nature first.

Modern technology is no longer a simple tool or extensions of the human body. The car is more than an expansion of the horse. The atom bomb is more than an extension of the sword. The sear magnitude of these projects has caused them to change their nature and have severed any connection with the past. Genetic engineering goes beyond a continuation of traditionally acceptable breeding practices. The size and scale of these endeavors put them in a qualitatively different category than earlier techniques. **Technology is not neutral, but ambivalent because it carries with it both good and bad effects quite apart from its use.**

Modern (New) Technology

Modern technology is unlike anything that has come before. Siegfried Giedion says that technology develops from a cultural milieu. "Tools and objects are outgrowths of fundamental attitudes to the world." According to Ellul, modern technology receives its preliminary beginning in the eighteenth century, but does not fully flourish until the second half of the nineteenth. Modern technology has a particularly unique existence that has no parallel in the past.

The techniques which result from applied science date from the Enlightenment and characterize our own civilization. Technique has taken substance; it has ceased to express a means as an intermediary. But is an object, an independent reality with which we must reckon.

Quantifiable explosion in contemporary technology has created a qualitative difference. This is the essence of the ontological analysis. **Technology is no longer adaptable to human society, rather society must adapt to technology.** Neutrality appears ridiculous in this light.

There is no value-free technology subject to the traditional limits of space, time and use, **"no longer conditioned by anything other than its own calculus of efficiency."** Technologists in the twenty-first century, realize the dangers of self-directed technology. Bill Joy says of recent scientific breakthroughs that they "pose a different threat than technologies that have come before. Specifically, **robots, engineered organisms, and nanobots share a dangerous amplifying factor: They can self-replicate.**" Modern technology behaves according to its own inner logic.

Characteristics of Pre-modern Technology

Local. The use and development of technology in the premodern world was generally limited to local areas. The development of a tool reflected the user's personality. Technology was part of civilization and culture; it did not constitute the whole and was not allowed to dominate. There was at best, a glacial progression of technology between cultures.

Choice. Prior to the advent of modern technology, choice was the decisive factor in what means an individual wanted to use for a particular task. Efficiency was only one choice. There was a greater diversity of choice in the tools one used prior to standardization.

The Ontology of Technology

Rationality and Artificiality

Rationality and Artificiality should be self-evident. The system functions according a strict internal logic and imposes a massive artificial megastructure over the earth, a technical environment that obliterates traditional cultures and ecology with a vast urban complex. A rational process over comes all spontaneous and emotional elements in human nature. There are divisions in labor, standardization and production norms spurns personal creativity, while rationalism follows a definite formula.

Automatism

Automatism represents the process in which technique will always follow, "the one best way."[212] This automatic procedure excludes personal choice. The technological imperative: what can be done must be done predominates. There can be no talk of limiting the most efficient means possible. We abdicate the right to choose in favor of efficiency. The decision made is merely following a mathematical law set in motion.

If the individual should oppose the operation of efficiency on moral or traditional grounds, he sets himself up against an enormous, "power against which there is no efficacious defense and before which he must suffer defeat."

Autonomy

Autonomy presents a social process that seeks to organize life around the principle of machine efficiency, "technique will assimilate everything to the machine; the ideal for which technique strives is the mechanization of everything it encounters." Technology creates an autonomous system that has broken the bounds of control and threatens to engulf culture and nature by transforming everything into a machine-like existence. "Technique . . . has fashioned an omnivorous world which obeys its own laws, and which has renounced all tradition."[215]

Self-Augmentation

Self-augmentation is the process where technology engenders itself assuring the continuity of its existence. An invention in one field leads to a multiplicity of applications in other fields - for example, the internal combustion engine.

Self-augmentation may be formulated into two laws. First, technical acceleration is irreversible. Technology never retreats or goes backward, unless society collapses. No return to an idyllic past is possible. There is no recall.

Second, "Technical progress tends to act not according to an arithmetic, but according to a geometric progression." An invention in one area multiplies itself in several others. This analysis is particularly alarming when applied to artificial intelligence and genetic engineering; for example, cloning will become at least as acceptable and common as in vitro fertilization; it will be another form of assisted reproduction and through the accumulation of genetic enhancements the human race will diverge into two separate species. One will be genetically modified, endowed with artificial intelligence, and superior to the other.

Two kinds of intelligence can never co-exist. One will rule the other one. This will mark the end of a distinctively human existence. The new genetically modified race will be considered posthuman. The modification of the human species will come through tiny individual changes in human life accumulated over many years.

Technological Risk

Risk is inherent to technological advance. Humanity is gambling on the idea that it can control what it creates. Think of the enormous power we hold in splitting the atom, splicing the gene, or cloning the embryo and in creating an artificial life. The argument that affirms given the right motivation technology can be put to good use fails at this point; **it does not consider the very real unintended consequences.**

Advanced technology will produce many benefits, but in the long run, it may eradicate the human species, albeit unintentionally. The balance of risk between potential benefits and cost analysis does not favor further acceleration, which has become irrational in light of the threat technology poses to life on earth. Technological risk endangers long term survival for short term gains and thus defeats the purpose of progress and human amelioration.

Idolatry

Technology is the great idol that faces down the human race and enslaves it to its own inner compulsion. Technology is the self-glorification of humanity to an unprecedented level; like its creator it is neither good nor bad both unreservedly both.

Technology in the Bible

Technology makes its first appearance according to Jacques Ellul in Genesis 4: 17-22. Here he declared, the sons of Cain established a city East of Eden with the ancestors who played instruments, such as the lyre and pipe and the forgers of bronze and iron tools. The resolve to defy God crystallizes in the Tower of Babel in the will to justify disobedience through construction of a city where technological power reigned supreme.

Theology of Technology: A Critical Approach

The *Brave New World* is not something in the remote future, but a process already in operation. Technical society drains the world of meaning and establishes an inhuman efficiency in its place. The pull of technicism into its vortex of meaninglessness by focusing on human dignity and choice must be rejected.

The Human Touch

We must seek to preserve the personal element in our witness. Berger remarked that "A major hazard of technology is its ability to provide deceptive substitutes for reality and relationships."[242] Professor Lanier Burns noted that people are, "most concerned about the loss of meaningful relationships, the human touch, in the mechanicalness of their technological world."[243] There must be a recovery of the importance of comradery. We must move away from mass evangelism and employment of media and marketing strategy and begin to think small, not in terms of reaching the masses, but in terms of reaching my neighbor, my family, my friends and co-workers. Evangelism should be thought of in terms of years and generations in the building of lasting friendships, and the raising of children, tailored to individual needs not crowds. Community is to replace networking. Emotionally coerced confessions of faith are not authentic. The gospel is not a commodity.[3]

[3] Terlizzese, L. J., Lawrence. (2019). *Killer computers: science fiction anticipates our future*. Cambridge, OH: Christian Publishing House.


POST-HUMAN
Genetic Enhancement &
Mechanical Augmentation
The Clouds Will Soon Roll By
Bradley W. Schrock, 2008


FRANCIS FUKUYAMA'S
"Our Post-Human Future"
The Clouds Will Soon Roll By
Bradley W. Schrock, 2008

During The 20th Century We Contained Existential Threat

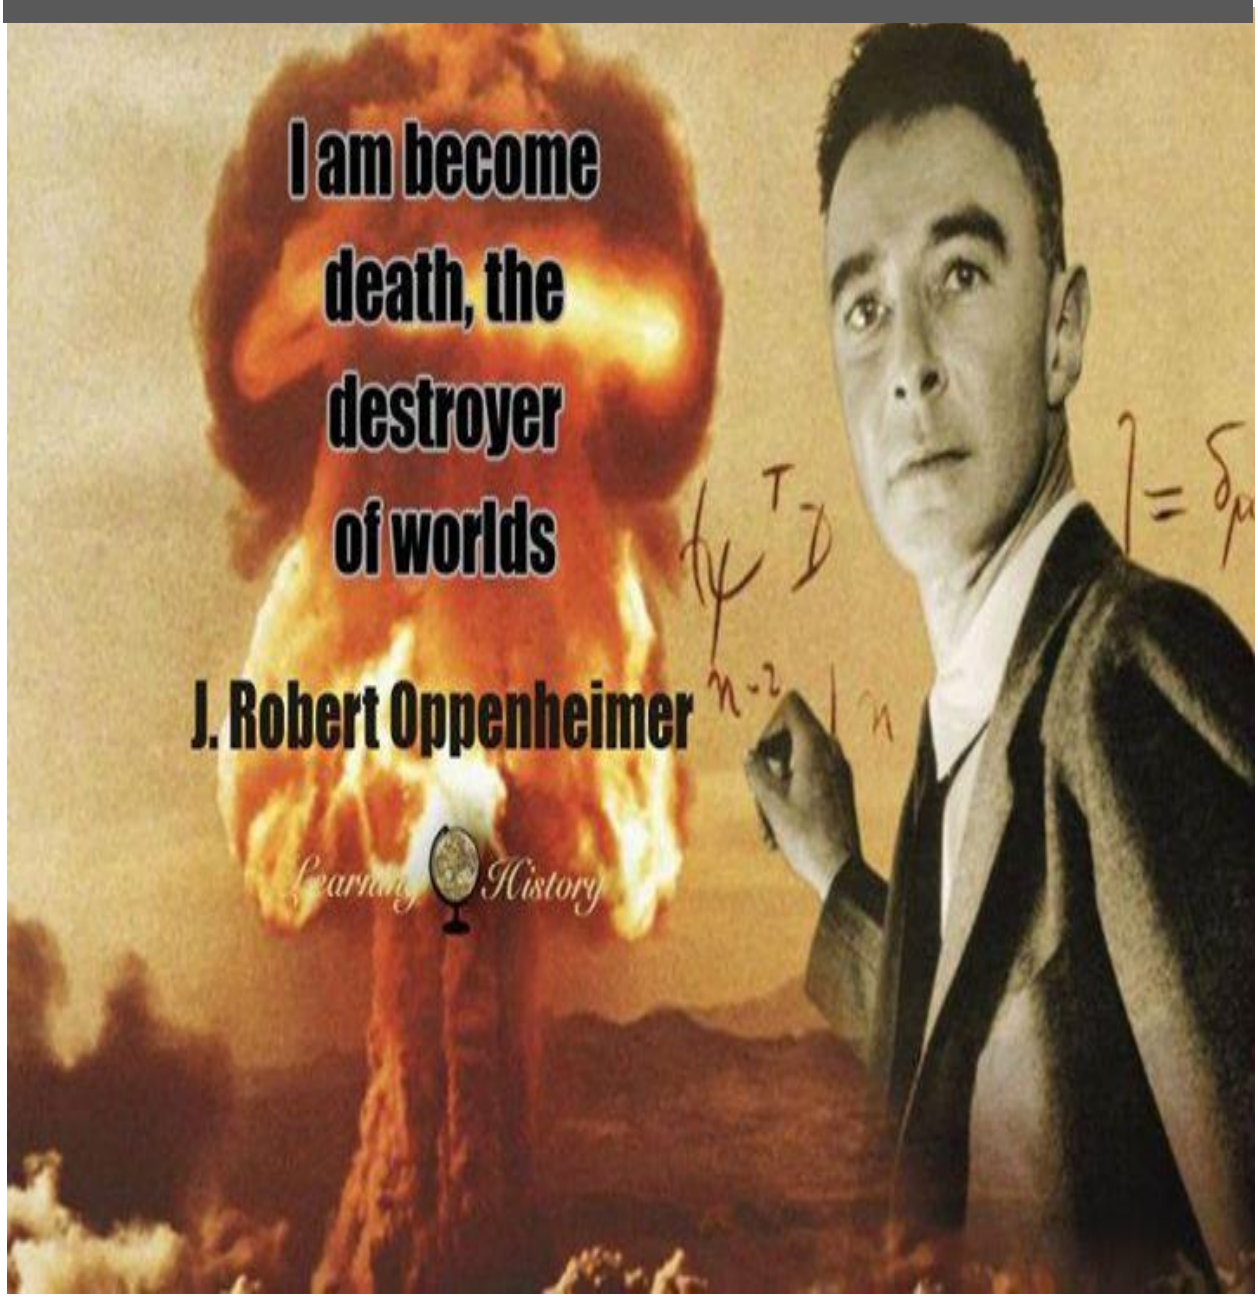

During The 21st Century We Can Tame The Threat Of AI

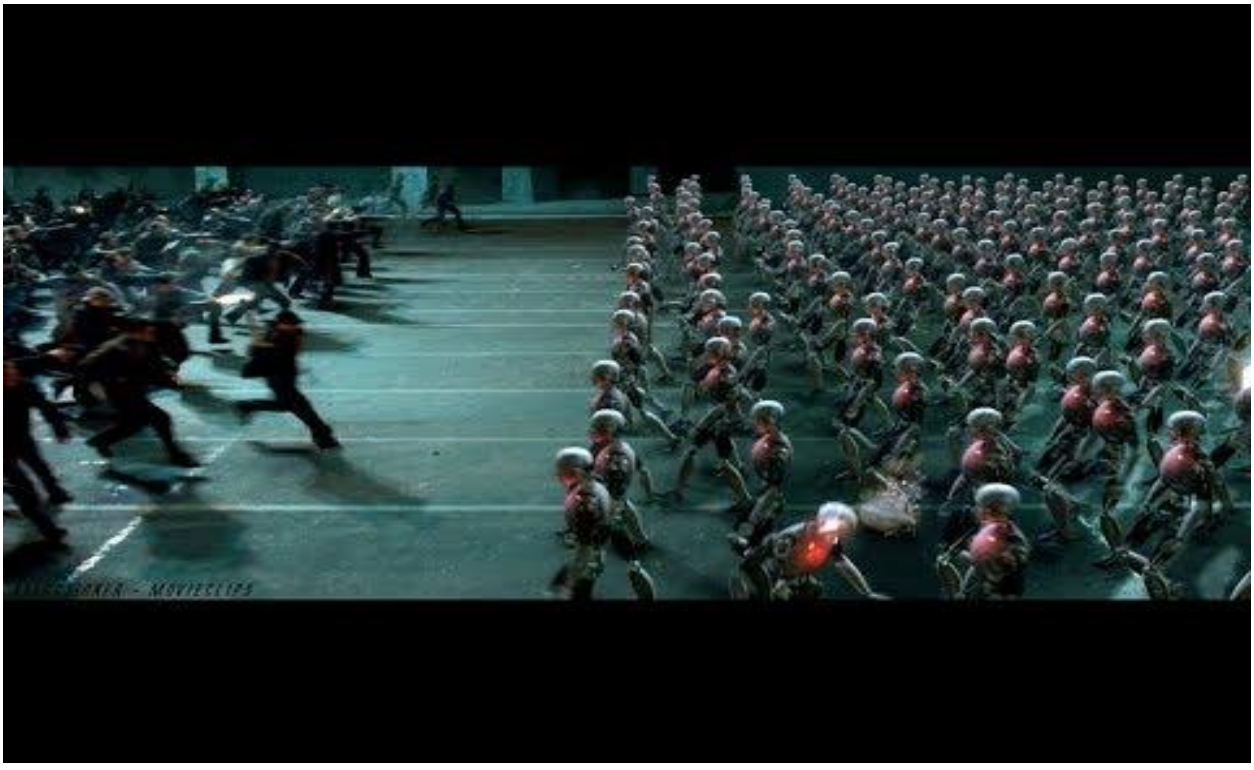

Concluding Summary & Personal Remarks

A recent podcast by Albert Mohler President of the SBC Flagship Theological Seminary seized my interest & initiated my research for this article. The subject of the podcast was Artificial Intelligence and Mohler's theme was the moral bias of Artificial Intelligence. Like with the proverbial bite into the Forbidden Fruit of the Tree of Knowledge of Good and Evil – warnings have been ignored – we will experience in unanticipated consequences for which we are not prepared. He refuted point by point claims for the moral neutrality of knowledge. The A.I. botnet is full of assumptions and beliefs. Chat applications aren't politically non-partisan. Bias is built into the cloud space and A.I. architecture making it an almost Herculean task to systemically remediate.

As the Texas Hold'em Poker Computer Challenge revealed – A.I. has demonstrated that it has reached a level no longer dependent on large volumes of data – only a set of rules. **It now teaches itself.** Bias mitigation controls will never be able to catch up until major damage has been done.

Combine this with the abdication of human decision-making, robotic replacement of the cognitive function in sensitive or critical professions/occupations and the disaster potential is greatly magnified. Do we face an exponential threat that can be described as an Existential Crisis? I would answer that we still have time to correct the situation if we get busy right now. Dystopian scenarios of Y2K did not take place because business and government cooperated to rapidly re-program.

Examples of Twentieth Century Science Fiction That Became Twenty First Century Science Fact Abound. While flip-top "communicators" - swiping/zooming – and even holograms have become present realities - much more of what was then predicted – especially by the eminent futurists at the top of their professions – never did occur – or at least not yet.

Personal Note: When I attended an Urban Planning Seminar while working on my Public Administration Master's Degree I was impressed by a futurist presentation of a speaker who represented the World Future Society. I was encouraged to join the group - which I did. New WFS members at that time were given a copy of a 1973 book produced by U.S. News & World Report entitled **"1994: The World Of Tomorrow."**

I still have my copy and after my 2023 re-read for this article - I am struck not by the few items they got right but by the majority they got wrong. Extrapolating from early seventies trends they confidently predicted human accomplishments such as - we would have people on nearby planets by 1980 – seven years after publication of the report. They predicted off-world colonization by the year 2000. In my opinion, the more predictably accurate were books by those unaffiliated with the WFS like *Future Shock by Alvin Toffler.*

Conclusion: Cold War Era Nightmare Dystopias Were Pro-Actively Prevented by Advancing Alternative Scenarios. In this Age of A.I. intervention must be focused on prevention of A.G.I. or Artificial General Intelligence especially without Safety Structure Systemically Governing Goal Architecture. To prevent rapid and efficient but soulless means-end output – there must be rules and limits – ethical constraints – even what's characterized as being moral boundary. – D. L. Burris