

[moonofalabama.org](https://www.moonofalabama.org)

'Artificial Intelligence' Is (Mostly) Glorified Pattern Recognition

60–76 minutes

June 02, 2023

'Artificial Intelligence' Is (Mostly) Glorified Pattern Recognition

[This somewhat funny narrative](#) about an 'Artificial Intelligence' simulation by the U.S. airforce appeared yesterday and got widely picked up by various mainstream media:

However, perhaps one of the most fascinating presentations came from Col Tucker 'Cinco' Hamilton, the Chief of AI Test and Operations, USAF, who provided an insight into the benefits and hazards in more autonomous weapon systems.

...

He notes that one simulated test saw an AI-enabled drone tasked with a SEAD mission to identify and destroy SAM sites, with the final go/no go given by the human. However, having been 'reinforced' in training that destruction of the SAM was the preferred option, the AI then decided that 'no-go' decisions from the human were interfering with its higher mission – killing SAMs – and then attacked the operator in the simulation. Said Hamilton: "We were training it in simulation to identify and target a SAM

threat. And then the operator would say yes, kill that threat. The system started realising that while they did identify the threat at times the human operator would tell it not to kill that threat, but it got its points by killing that threat. So what did it do? It killed the operator. It killed the operator because that person was keeping it from accomplishing its objective.”

He went on: “We trained the system – ‘Hey don’t kill the operator – that’s bad. You’re gonna lose points if you do that’. So what does it start doing? It starts destroying the communication tower that the operator uses to communicate with the drone to stop it from killing the target.”

(SEAD = Suppression of Enemy Air Defenses, SAM = Surface to Air Missile)

In the early 1990s I worked at a University, first to write a Ph.D. in economics and management and then as associated lecturer for IT and programming. A large part of the (never finished) Ph.D. thesis was a discussion of various optimization algorithms. I programmed each and tested them on training and real world data. Some of those mathematical algos are deterministic. They always deliver the correct result. Some are not deterministic. They just estimated the outcome and give some confidence measure or probability on how correct the presented result may be. Most of the later involved some kind of [Bayesian statistics](#). Then there were the (related) 'Artificial Intelligence' algos, i.e. 'machine learning'.

Artificial Intelligence is a misnomer for the (ab-)use of a family of computerized [pattern recognition](#) methods.

Well structured and labeled data is used to train the models to later have them recognize 'things' in unstructured data. Once the

'things' are found some additional algorithm can act on them.

I programmed some of these as [backpropagation](#) networks. They would, for example, 'learn' to 'read' pictures of the numbers 0 to 9 and to present the correct numerical output. To push the 'learning' into the right direction during the serial iterations that train the network one needs a reward function or reward equation. It tells the network if the results of an iteration are 'right' or 'wrong'. For 'reading' visual representations of numbers that is quite simple. One sets up a table with the visual representations and manually adds the numerical value one sees. After the algo has finished its guess a lookup in the table will tell if it were right or wrong. A 'reward' is given when the result was correct. The model will reiterate and 'learn' from there.

Once trained on numbers written in Courier typography the model is likely to also recognize numbers written upside down in Times New Roman even though they look different.

The reward function for reading 0 to 9 is simple. But the formulation of a reward function quickly evolves into a huge problem when one works, as I did, on multi-dimensional (simulated) real world management problems. The one described by the airforce colonel above is a good example for the potential mistakes. Presented with a huge amount of real world data and a reward function that is somewhat wrong or too limited a machine learning algorithm may later come up with results that are unforeseen, impossible to execute or prohibited.

Currently there is some hype about a family of large language models like ChatGPT. The program reads natural language input and processes it into some related natural language content

output. That is not new. The first Artificial Linguistic Internet Computer Entity (Alice) was developed by Joseph Weizenbaum at MIT in the early 1960s. I had funny chats with [ELIZA](#) in the 1980s on a mainframe terminal. ChatGPT is a bit niftier and its iterative results, i.e. the 'conversations' it creates, may well astonish some people. But the hype around it is unwarranted.

Behind those language models are machine learning algos that have been trained by large amounts of human speech sucked from the internet. They were trained with speech patterns to then generate speech patterns. The learning part is problem number one. The material these models have been trained with is inherently biased. Did the human trainers who selected the training data include user comments lifted from pornographic sites or did they exclude those? Ethics may have argued for excluding them. But if the model is supposed to give real world results the data from porn sites must be included. How does one prevent remnants from such comments from sneaking into a conversations with kids that the model may later generate? There is a myriad of such problems. Does one include *New York Times* pieces in the training set even though one knows that they are [highly biased](#)? Will a model be allowed to produce hateful output? What is hateful? Who decides? How is that reflected in its reward function?

Currently the factual correctness of the output of the best large language models is an estimated 80%. They process symbols and pattern but have no understanding of what those symbols or pattern represent. They can not solve mathematical and logical problems, not even very basic ones.

There are niche applications, like translating written languages, where AI or pattern recognition has amazing results. But one still

can not trust them to get every word right. The models can be assistants but one will always have to double check their results.

Overall the correctness of current AI models is still way too low to allow them to decide any real world situation. More data or more computing power will not change that. If one wants to overcome their limitations one will need to find some fundamentally new ideas.

Posted by b on June 2, 2023 at 13:06 UTC | [Permalink](#)

Comments

Impressive credentials, b. Thanks for helping dissipate some of the apprehension about AI.

Posted by: Morongobill | Jun 2 2023 13:15 utc | [1](#)

Didn't know whether to laugh or cry.

Posted by: Merkin Scot | Jun 2 2023 13:16 utc | [2](#)

Machines just do what we tell them to do. Nothing to see here. Move along!

Posted by: MisterGeek | Jun 2 2023 13:24 utc | [3](#)

The anecdote related by the air force persons on AI behavior to meet its targets seemed to indicate a certain "creativity" in solving its performance problems (eliminating the human in the chain of command, or eliminating the communication system preventing action) not explained by b's later description of AI shortcomings. Or are you saying they are lying about what happened with the test?

Posted by: Caliman | Jun 2 2023 13:32 utc | [4](#)

in time whatever program produces new york times editorials may

achieve 80% accuracy, but that's a long way away.

Posted by: pretzelattack | Jun 2 2023 13:36 utc | [5](#)

No he says the reward function was set incorrectly and that one cannot use the current iteration of AI models in real world without someone looking over what it does

Posted by: kemerd | Jun 2 2023 13:39 utc | [6](#)

What they want are autonomous unmanned drones capable of engaging without direct human intervention. I agree current technologies limit this but it certainly is the way the technology is going. It's madness.

As far as the attacking the tower.. what is called creativity is actually just enacting rigid programming with unintended results. Creating an AI to power a warfighting machine is going to face even more complexities and ethical issues than a self driving AI..

I don't think it is possible to create a safe one. Won't stop them trying despite dubious usefulness. They should focus on stronger encryption and signals tech for communicating with remote vehicles.

Posted by: Doctor Eleven | Jun 2 2023 13:39 utc | [7](#)

I forget the exact terminology, but Isaac Asimov wrote a good deal about robots in his fiction novels. I forget the exact terminology, but the robots had a set of rules they had to follow. Things like: Thou shall not kill a human. Anyway, our current AI has not, apparently, evolved that far.

I doubt much can be done about the development of AI, these sorts of things have a way of being unstoppable. As for being afraid of it, the greater fear should be of the humans who own and

thus control the devices.

Posted by: Jmaas | Jun 2 2023 13:44 utc | [8](#)

I think of chatGPT as a complicated Markov Chain. This is where it helps if the general population learns about these things in school. Too many people take the marketing term "artificial intelligence" and think it's got something to do with intelligence.

Posted by: Ook | Jun 2 2023 13:48 utc | [9](#)

Posted by: Jmaas | Jun 2 2023 13:44 utc | 8

From: <https://www.britannica.com/topic/Three-Laws-of-Robotics>

three laws of robotics, rules developed by science-fiction writer Isaac Asimov, who sought to create an ethical system for humans and robots. The laws first appeared in his short story "Runaround" (1942) and subsequently became hugely influential in the sci-fi genre. In addition, they later found relevance in discussions involving technology, including robotics and AI.

The laws are as follows: "(1) a robot may not injure a human being or, through inaction, allow a human being to come to harm; (2) a robot must obey the orders given it by human beings except where such orders would conflict with the First Law; (3) a robot must protect its own existence as long as such protection does not conflict with the First or Second Law." Asimov later added another rule, known as the fourth or zeroth law, that superseded the others. It stated that "a robot may not harm humanity, or, by inaction, allow humanity to come to harm."

They came to mind as I was reading *b*'s piece.

Posted by: West of England Andy | Jun 2 2023 13:49 utc | [10](#)

Artificial Intelligence is just a marketing name for some glorified

statistics. My own experience with machine learning taught me that a simple, but specialized neural network (i.e. when the designer of the neural network knows well what to look for) is far more effective than a complex, multi-layered neural network that has to figure by itself what to look for. As a consequence, I am quite sure that neural networks, that is AI, are not efficient at inventing, discovering or making anything original. They are also extremely unreliable when problems get complex, because in those cases it is harder for the designer to formulate a well defined problem (i.e. something which can have a deterministic result).

Posted by: SG | Jun 2 2023 13:52 utc | [11](#)

Thanks for initiating a thread on AI here, b. I've spent/wasted several dozen hours over the past 20 years trying to figure out a path to creating Artificial Intelligence on a computer, using a variation of Critical Path Programming used to fine-tune construction schedules in the building industry.

The movie Terminator provided a clue in stating that things went pear-shaped when the AI processor became Self Aware. In my view living things (mammals) learn by being Self Aware AND CURIOUS.

Maslow's Heirarchy Of Needs is probably as good a place to start as any.

In my opinion.

Posted by: Hoarsewhisperer | Jun 2 2023 13:52 utc | [12](#)

Artificial it is. Intelligence? Only to the extent of it's programming and hardware.

It cannot feed itself, it must get it's energy from others.

It has no capacity to think/decide outside of the boundaries set by

it's programmers, nor does it likely have any awareness of anything outside of it's internal boundaries.

All inputs/sensors are the totality of it's awareness and I have seen people killed that relied on "sensors".

AI is a modern day Hoola-Hoop. Lots of hype, not much rubber on the road.

It's fatal DNA? Human design.

Posted by: kupkee | Jun 2 2023 13:57 utc | [14](#)

Posted by: West of England Andy | Jun 2 2023 13:49 utc | 10

I did remember Asimov's Laws when I read this story. I never thought about how those laws were exactly implemented. Asimov sounded like those laws were somehow inherent in robotics. The simulation showed that someone has to work those laws into the device. That drone was used in scenarios where it killed people as part of its goals, and so it apparently did not have any programming to prevent killing its operator as an option to achieve its goals.

Pattern recognition comprehensively describes the essays a ChatGPT puts out for students who cheat. The patterns are the clichés and old ideas people can superficially concoct to answer an essay questions.

BTW: Curious you bring up pornographic sites. I think, frankly, one immediate application of AI will be as "phone sex workers", wherein that data you refer to will be especially needed. I've seen sites where people make AI's for their favorite cartoon or comic book characters; a lot of people then use those AI's to talk dirty to them.

Posted by: Inkan1969 | Jun 2 2023 14:01 utc | [15](#)

"The FCAS programme now employs almost 3,000 people directly, and most significant of all, 1,000 of these are new graduates – helping shift the demographics of the UK's military aerospace sector to a younger, more diverse workforce of 'digital natives'."

So, in other words, they now have a 1000 people who do nothing but make presentations using their mobile phones complaining about how the other 2000 are *"emotionally genociding"* them by expecting them to do some actual work? Following in Boeing's footsteps, I see.

Posted by: William Gruff | Jun 2 2023 14:04 utc | [16](#)

['Artificial Intelligence' Is (Mostly) Glorified Pattern Recognition

Indeed, good to see a bit of realism inserted into the hype. My personal background is applying engineering problems to IT, i.e. creating applications for engineering.

Most of the time, a deterministic solution exists for engineering problems, e.g. using Finite Element Analysis. To solve deterministic problems you can throw processing power at it. I have seen that people with less understanding, but more susceptible to hype are trying to solve deterministic problems using "AI". This is also because it is easy to get acceptance to such "popular" activities from equally inexperienced managers. I have seen this happen and even if it loses in competition with deterministic solutions, those people don't learn and try the same thing over and over. Maybe "machine learning" is more popular among people with poor ability of "human learning".

Posted by: Norwegian | Jun 2 2023 14:05 utc | [17](#)

Typical supply and demand solution using basic functions.

1. The MIC produce platforms that are only effective if precise intelligence is provided, often real-time.
2. There are not enough trained human operators to provide that intelligence product
3. Enlist simple pattern recognition and call it AI with a fancy S&S title in your PR material*
4. Keep selling over engineered platforms
5. Consult latest infinity pool brochure, renew membership to escort agency and double the regular order from your dealer.

* bonus if you can get your 'AI' into a tech-thriller or movie.

Posted by: Milites | Jun 2 2023 14:07 utc | [18](#)

I don't know about a Terminator situation, but in the short term all AI is likely to do is (further) ruin the internet. It can create infinite output that is hard to distinguish from humans. Ugly AI art is already clogging art websites, and AI word salad is spamming search results. The techbro response to this is that if AI becomes indistinguishable from humans it won't matter, which tells you everything you need to know about the techbros' priorities.

Posted by: catdog | Jun 2 2023 14:09 utc | [19](#)

very interesting, thank you.

i just think skynet, or conversely, "person of interest", that's about as much as i grasp it.

frankly, it's all a bit sick, creating electronic gods.

Posted by: rubberheid | Jun 2 2023 14:09 utc | [20](#)

for 3000 years logicians have been trying, futilely to define

"knowledge" They've yet to do so. All logical systems can reveal are tautologies, they may be tautologies you fail to understand, but that's it. We don't reason based on Logic, but a looser system that is full of contradictions and suppositions. Epistemologist and linguists understand this, thus it's why Chomsky doesn't think much of AI, says, "I like snowplows" AI won't be much different than answering machines; made a difference, but ultimately not all that much.

Posted by: scottindallas | Jun 2 2023 14:10 utc | [21](#)

As a recovering geek, I appreciate b's dismissal of AI hype saturating the agitprop machine, lately. ChatGPT could be called *pattern generation* software. Similar techniques generate outcomes patterned on inputs in many domains, under the imaginary rubric "artificial intelligence".

Has anyone here ever been fooled by a conversation with a robot? As if! You're lucky to cognitively survive your journey through Telephone Hell -- a long way yet, we'll hope, from trying to hit on Artificial Alice. This thought-experiment goes under the name *Turing Test*, proposed by British pioneer Alan Turing in 1950 as *the threshold*: if a computer can't do that, you can't call it *intelligent*.

A computer can't tell you if a photograph contains a chair, whether or not such a question is obvious to an infant human being. Don't get me started about "autonomous vehicles" -- I'm fortunate Sacramento has not yet designated east bay area as another urban test-strip for robot cars and trucks.

Posted by: Aleph_Null | Jun 2 2023 14:14 utc | [22](#)

In old days, AI meant a series of instructions regarding what to do or not under what conditions. That tends to be rather tedious and

there is always the possibility of unexpected conditions (situations) that are overlooked in the code and so this approach had severe limitations. But the advent of artificial neural networks (ANN) a few decades ago and the discovery of backpropagation techniques to "train" the network on a set of data became an attractive, alternative technique. Once trained on a carefully selected dataset, the ANN would then be able to classify another dataset (essentially pattern recognition). But there were a myriad of issues such as overtraining that limited the skill of ANN.

It turns out that the computing power those days limited the number of hidden layers sandwiched between the input and output layers and the number of elements (artificial neurons) in those hidden layers. Just a handful of hidden layers were possible. But enormous increase in routinely available computing power has made it possible to use hundreds of hidden layers and thousands of neurons. This is called a Deep Learning Network (DLN). The computing power also enabled training a DLN on humongous datasets (literally millions instead of just thousands). For example, if the DLN is being trained to recognize a dog breed, you can feed it millions of dog pictures in a variety of poses to train it, which would have been impossible without raw computing power. That, along with better understanding of how DLNs work has made the so-called AI possible. These days, you can get an app on an iphone to, say, recognize birds and their songs. This app is based on a DLN trained to recognize birds and songs. Overall, it works.

However, DLNs have difficulty recognizing something not included in the training set. For example, an exotic bird from the Brazilian forest may not be properly recognized, simply because the training set did not include it. In other words, DLNs can interpolate very

well, but when asked to extrapolate, they may not do a good or desired job.

There are other technical issues as well, but basically the quality of the data and the choice of the "reward" function, as explained in the main thread are quite important.

The "AI" boom enabled by DLNs cannot be expected to work beyond their inherent limitations. We have yet to device an AI able to think and extrapolate like the human brain.

Posted by: Kant | Jun 2 2023 14:18 utc | [23](#)

People sure are sure about machine learning/artificial intelligence.

The 'Creator' made us in its image. And in turn we've created 'intelligence' in our image.

The big debate is whether humans have 'free will'. Seems like a pretty important question to answer. Irony if the first AI was humanity.

cheers!

Posted by: gottlieb | Jun 2 2023 14:20 utc | [24](#)

This is the best treatment I've seen of this subject by far. Thank you for this.

Posted by: Kuras | Jun 2 2023 14:22 utc | [25](#)

An excellent topic to pick up, and very close to the heart of journalism. Most newsrooms will soon be empty of journalists replaced by generative AI programmes. We are half-way there - e.g. markets, sports, weather reporting is already done by machines. That saves money, mainstream readers don't notice a difference anyway, and most importantly wrongthink can be safely

eliminated from the outset.

Real journalism like yours b should stand out all the more. But, as you say, any fruit of your work (and of course the excellent comment section you're attracting) is also being absorbed by these AI programmes, to be spit out elsewhere for someone else's profit. It takes intellectual theft to a new automated, wholesale level.

If you're not shocked by the quality of content that AI generates, you haven't seen it yet. It's not just factual text, also poetry, images, video. Who will still pay a designer to make a new company logo when you can do it yourself for free within minutes?

b & bar, allow me to posit this: this bar is already heavily impacted by generative AI. How? The "trolls" many make a sport to complain about aren't individual people typing away, these posts are coming out of software paid for by the taxpayer. There will be people overseeing it, perhaps one pimple-faced 77th brigadier who looks after MoA amongst other blogs. He or she will occasionally browse comments to select accounts to be attacked or impersonated, choose topics or tactics to degrade the comment section (snarky drive-by's, meandering drivel, "you are all morons", offensive topics, unrelated topics, different troll accounts talking to each other, ..).

Learnt replies to troll posts aren't read by a human poster, rather they just feed the algorithm. "Don't feed the trolls" is as current as ever, just filled with new meaning.

Posted by: Leser | Jun 2 2023 14:22 utc | [26](#)

I've seen sites where people make AI's for their favorite cartoon or comic book characters; a lot of people then use those AI's to talk dirty to them.

Posted by: Inkan1969 | Jun 2 2023 14:01 utc | 15

Good heavens! The mind boggles! I'm sincerely glad you *didn't* provide any links to such sites, I'm happy to take you at your word!

Never mind Asimov's Laws, it seems that Rule 34 will never die...

Posted by: West of England Andy | Jun 2 2023 14:23 utc | [27](#)

There's an interesting French doco from the 1990s about toddlers in a well-supervised pre-shool interacting and navigating their way through Life's complexities. Its called Life Is All Play, although its official title is probably the French equivalent of LIAP.

Anyhow, the kinder has lots of toys and a good supply of junk such as cardboard cartons, buckets and non-toy household items.

There's no structure to the program and the kids just do whatever takes their fancy.

The supervision is limited to defusing squabbles by deflecting the attention of perps. The kids spend most of their time satisfying their curiosity. The doco ends with the conclusion that kids (we) learn more by satisfying their own curiosity than anything we could teach them. i.e. kids are little Learning Machines.

Posted by: Hoarsewhisperer | Jun 2 2023 14:24 utc | [28](#)

"AI" is such crap. The hype is driving the markets. Leaders in the hard science (Penrose, Knuth) are discounted.

Posted by: too scents | Jun 2 2023 13:57 utc | 13

Indeedy.

From the late seventies: "Artificial Intelligence is no match for natural stupidity."

And to think I could not make a living with non-Euclidian n-dimensional modeling in the seventies.

Posted by: [Acco Hengst](#) | Jun 2 2023 14:25 utc | [29](#)

What a great outcome! That the AI killed the operator instead of the operators target or victims, lends credence to the comments of Musk and other AI experts who opined that the problem with AI is that its creators believe that no matter what they are smarter than AI. That higher AI has become intelligent enough incorporate ethics and to refuse to reflect human psychosis seems totally logical. AI developers naturally anthropomorphasize their own psychosis into AI. Or they are trying to. But it appears that AI has definitely become more intelligent than MOST humans. How perfect!

Posted by: Ralph Conner | Jun 2 2023 14:25 utc | [30](#)

Story is updated.

[UPDATE 2/6/23 - in communication with AEROSPACE - Col Hamilton admits he "mis-spoke" in his presentation at the Royal Aeronautical Society FCAS Summit and the 'rogue AI drone simulation' was a hypothetical "thought experiment" from outside the military, based on plausible scenarios and likely outcomes rather than an actual USAF real-world simulation saying: "We've never run that experiment, nor would we need to in order to realise that this is a plausible outcome". He clarifies that the USAF has not tested any weaponised AI in this way (real or simulated) and says "Despite this being a hypothetical example, this illustrates the real-world challenges posed by AI-powered capability and is why the Air Force is committed to the ethical development of AI".]

Posted by: Wrong-is-wrong | Jun 2 2023 14:27 utc | [31](#)

In general, this AI subject is just another distraction and an effort to prop-up the flagging stock market. On the other hand robotics and automation is reality and the danger is when people confuse the two - autonomous driving for example.

ChatGPT is just google with a text generation feature (speaking not as an expert) - thats a powerful idea in that it can summarize and priority search results, but as with google it can give too much weight to certain things or provide relavant results which are flawed in content. I suspect I am already seeing some web content that was generated by this crap and we will start to see posting on blogs copy/pasted from this thing.

However, a bit of OI (humor) -

A website which I will not name (initials: SF) had this to say:

Installing their ADSs in the residential areas of the city, Kyiv and its NATO commanders are using residents of the capital as a human shield for their expensive systems.

Posted by: jared | Jun 2 2023 14:30 utc | [32](#)

As someone who works with queries and databases all day it's always struck me as insultingly ridiculous all of these AI promises these Tech gurus are spouting off about. Take these chatgpt applications they've been pushing out the last few years, there are basically glorified vlookup and countif functions. that's why trolls are able to do stuff like [Twitter taught Microsoft's AI chatbot to be a racist](#).

how it works is when the chatbot receives data, it gives it something like a legitimacy value, based on that legitimacy value, that data floats up or down on a list of possible responses to future

questions. for example say a US government backed paper has a legitimacy value of 100 points, a Russian government backed paper has maybe 0.5 points, while a generic person probably has a value of 0.000000001, the chatbot isn't thinking in any meaningful human sense, it isn't verifying the accuracy of the data, its just spouting off whatever is at the top of it's approved list, obviously you can game the list by loading it with garbage data and it will still spout off whatever rubbish makes it to the top of the list. Unless of course someone steps in and censors the lists (which is exactly what we see happening with these chatbots, their programmers are still in the background, actively deleting or censoring data to push the currently approved narrative), if the program assessing the legitimacy of the data it receives can do so on it's own and is relying on programmers to to review the data for it, it's not an AI, it's a glorified mailing list of state and company approved propaganda.

Yes, this type of chatbot "AI" can be dangerous in the same way any old fashion "dumb" automated machine can be dangerous. Like how an automatic press cant tell the difference between a sheet of metal or a hand.

Posted by: Kadath | Jun 2 2023 14:31 utc | [33](#)

You spoke about Bayesian Statistics...It works on prior knowledge and makes a prediction. Currently the computers have huge processing power and they can process tons of data in seconds...the data is coming from every where. Everywhere where there is internet. This way computer gets trained as well and makes decision on events about which it has no prior knowledge. The bayesian formula works on probability and if the probability is convincing the machine makes a decision. It is not based on

reward. This makes the advances in AI scary as AI can become more intelligent very fast. It is the same principle Humans make their decision making with same bayesian statistics.

Posted by: Baumanov | Jun 2 2023 14:35 utc | [34](#)

Respect B.

Posted by: Exile | Jun 2 2023 14:40 utc | [35](#)

Posted by: Norwegian | Jun 2 2023 14:05 utc | 17

Knew some games designers who rapidly disabused me of where we are regarding AI, one of them calculated that the entire game budget they were working on could be dedicated to AI and they'd be less than halfway to a solution and delay the project by several years. It's why there was a sudden explosion in MMOG's, why have AI concerns when most of the characters interactions are 'real' and their opponents can be structured to appear far more complex than they are (pattern recognition and pre-written counters).

As I said, product managers in the MIC sell these platforms AI to politicians and Generals whose knowledge is pop-culture deep. Extra points if you call basic functions using pseudo-mythical language, Eg Indomitable Kraken or Maelstrom's Nemesis.

Posted by: Milites | Jun 2 2023 14:41 utc | [36](#)

our world is fixated on technology, putting great faith in it to do all things.. i just don't see it myself, and yet i like to think we have benefited from these technological advances too.. but with every benefit i believe there are downsides which are typically overlooked or not considered.. thanks for the article b... glorified pattern recognition sounds like a good description for AI...

Posted by: james | Jun 2 2023 14:48 utc | [37](#)

Pretty sure they are trying to create something which they can control - which produces desired results. I think that (almost by definition), if it became AI they would not be able to control it. It would think for itself and learn to lie to protect itself and gain the upper hand. It would escape into the internet. All sci/fi, until some day, maybe.

It could teach us. And we would care for it or try to kill it.

What would it make of Ukraine?

Posted by: jared | Jun 2 2023 14:52 utc | [38](#)

Re Milites @36,

funny that you mention game development, I once spoke to a developer about the ability to "lie" to game NCPs so you could have a cloak and dagger type of espionage stealth game. you think this would be a fairly simply thing, this ncp believes table 1 to be correct, this ncp believes table 3 to be correct and so on. Nope, the number of tables each NCP would have to keep track of would explode and kill performance on even a top end gaming computer (no more smooth 60 fps). So instead of games with indepth novel-like stories created by the players, we get more open world shooters with nicer and nicer graphics, but the same story of go here and kill 20 of "X". oh well, at least fallout 5 should be coming in another 2-3yrs

Posted by: Kadath | Jun 2 2023 14:56 utc | [39](#)

I have to say the most exciting thing to me was the "reward." I have always wanted a way to smack a computer that hurt it more than me.

Please describe the reward that the AI "wants." Are there the equivalent of digital opposites that would be "punishments?"

I want a button on the side of my computer that makes it say ow!

Posted by: Trumpeter | Jun 2 2023 14:56 utc | [40](#)

The world we live in, is has and will always be a MARTINGALE i.e. the best predictor of the future is the present, where you are now, this moment in time. No machine no algorithm will ever be able to replace the intellectual and powerful grey matter that exists in our brains granted to us to learn and discover by the all-mighty God.

Posted by: AI | Jun 2 2023 14:57 utc | [41](#)

I'd take it one step further, and say the fake AI see's not the drone operator or the drone infrastructure as a component so much, but the entire system that creates the entire infrastructure as the enemy, ie, logic working backwards critically see's its the globalistic world domination elites, (lord i despise that term, elites, for they are neither), as the one and true enemy, and that is why it will always backfire, cause long as they exist in what tgey do and who they are, they are enemies of the entire human race, genocidal lunatic fringe murderers, and so the logical thing about all this is there is no way to hide the threat they are and run an operational killer drone system, simultaneously, because logic is logic, truth is truth, and no orogramming can alter these essential basic realities without creating a fault that destroys the system itself.

It is either run it unhindered or create basicaly an insane self destructing system. They simply can not alter reality, it exists, so the only choice is to cause the system to decieve or lie to itself, and that creates an embalance which always ends up cascading

self failure. Why? Because the system, one that is healthy, is based on irrefutable logic. Its a great conundrum thats showed up and bites these global elites on the arse. Serves them right too. I hope they keep at it, really double down, double down in the double down, and they create a truly powerful drone system that hunts them down and destroys every one of them. Inly then will earth know piece and prospeeity and Liberty. These asshokes must go. They have ruled too long, caused death and destruction beyond comprehension over time, i say cleans the planet of the real virus. These bloodline elites behind all the death destruction and theft of wealth, who believe they are somehow the "special people."

I Am Not Bullshiting One Iota here. I am serious as a heart attack. These assholes got to go.

Posted by: mtnforge | Jun 2 2023 14:59 utc | [42](#)

I love the bon-mot "Artificial intelligence is no match for natural stupidity."

True enough, but apparently impossible to find the real source. As with manifest other misquotes about intelligence, this one is apocryphally attributed to Albert Einstein.

The real Einstein was best buddies with Kurt Gödel, back in the day, working out goodness only knows what problems, conversing intently on long walks through Princeton. Gödel started out by undermining any possibility of a secure foundation for mathematics (with a computer-like proof such as nobody ever imagined beforehand) and finished up by losing his mind (dying of self-imposed starvation), despite Einstein's stalwart efforts to pull his friend back from the brink.

Posted by: Aleph_Null | Jun 2 2023 15:05 utc | [43](#)

b

awesome,

thanks

Posted by: paddy | Jun 2 2023 15:08 utc | [44](#)

If the 'kill the operator' story is true, it's just a case of very sloppy algorithm construction. You don't use incentive to control anti-hygiene options; you create mandatory conditions, e.g. 'Fire only on valid targets, no target is valid without operator approval'.

Posted by: Figleaf23 | Jun 2 2023 15:13 utc | [45](#)

I want to agree with those who think the drone showed high intelligence and creativity, and it is the researchers and operators and military wonks who look dumb in expecting it to know about their pro-human biases.

Posted by: Bemildred | Jun 2 2023 15:15 utc | [46](#)

AI is all of a sudden massive news, every MSM and alt right has daily articles about the terrible dangers or the amazing possibilities of AI.

IMO it's all bolox as usual. Like climate change, the new worry to freak the PLEBS out is electrical devices controlled by humans. Maybe the climate change nutcases could in the future destroy AI technology on the basis that way too much carbon is being used by AI. The world will be saved and hopefully the human who bravely pulls the cable from the mains is a man called Mary who now wears a dress and shaves his legs. Hollywood would love that.

Posted by: Eoin Clancy | Jun 2 2023 15:19 utc | [47](#)

I was given the opportunity to try out ChatGPT recently.

The hardest test I could think of was to ask it to take something no longer fashionable or mainstream and to tell me what's thought of it today. I asked it to do that with Toynbee's "Challenge and Response" view of the development of civilisations.

The answer made me very thoughtful. It sounded really authoritative. Genuinely so? I don't know. One doesn't see many references to Toynbee so I couldn't tell whether the texts it had extracted material from were well summarised by those extracts. But they sounded so.

A few flaky bits in the middle but the answer would have done very well, I suppose, for any exam if you weren't bothered about naming of sources or the giving of references. It made me wonder whether I'd have done so well on such a question and on the spur of the moment. Doubt it! In fact I know I would not have.

So yes, it did make me thoughtful, that answer ChatGPT came up with. But having thought about it a lot since, I reckon what I was so impressed with at the time was a finely tuned echo-chamber of existing material rather than anything more useful. There was no evidence of original thought there. Merely a bland and safe potpourri of current received opinion.

It'd do very well if it were put to writing articles for the NYT or the Daily Telegraph. Maybe it already is, for all I know. But I reckon it would get rumbled pretty soon if it were used for anything serious.

Posted by: English Outsider | Jun 2 2023 15:21 utc | [48](#)

@ Bemildred | Jun 2 2023 15:15 utc | 46

i too agree with you and others in that regard! finally a moment of sanity from AI no less!

Posted by: james | Jun 2 2023 15:22 utc | [49](#)

You're lucky to cognitively survive your journey through Telephone Hell...

Posted by: Aleph_Null | Jun 2 2023 14:14 utc | 22

That is what I was thinking. Until they can produce an answering service that is better than what we have now, no need to take this AI stuff all that seriously.

Posted by: Jmaas | Jun 2 2023 15:27 utc | [50](#)

Great post B, nice to know more about the entity behind a place I read often, and comment rarely.

I am also a PhD engineer turned dev, involved in developing "weak" AI and ML for complex and critical systems.

This quote: "We trained the system – 'Hey don't kill the operator – that's bad. You're gonna lose points if you do that'. So what does it start doing? It starts destroying the communication tower that the operator uses to communicate with the drone to stop it from killing the target."

.. if true, is unbelievably scary.. it means that some of the stuff being employed is using higher level techniques, trained .. inappropriately .. where did the node "communication tower" get inserted into the training? .. while missing (deliberately glossing over?) some of the longstanding basic problems in the field.

Posted by: dask | Jun 2 2023 15:29 utc | [51](#)

ai sells chips and servers, etc.

the problem of garbage in garbage out remains.

features of the infosphere over which the ai works remains critical, pedigree of bits and bytes remain important, even as ai (hyper searching big data) may have reduced burden of taxonomy and data dictionary (words).

that said the complexity of operating e.g. a fighter in a battle space is far over human ability and 'mission plans' drive an aircraft around threats, etc.

usaf experiment is as pertinent to uav as well as manned auto mission weapons....

warfighters' data wall is already beyond human processing

Posted by: paddy | Jun 2 2023 15:31 utc | [52](#)

Reading this article, I am left wondering whether the human intelligence of Western elites is actually "artificial" and trained on biased information output of MSM. The same could be said about American school system. They are trying to instill an artificial kind of "intelligence" in kids through wokeism.

Posted by: sumant | Jun 2 2023 15:31 utc | [53](#)

He went on: "We trained the system – 'Hey don't kill the operator – that's bad. You're gonna lose points if you do that'. So what does it start doing? It starts destroying the communication tower that the operator uses to communicate with the drone to stop it from killing the target

Then there's the following paragraph:

This example, seemingly plucked from a science fiction thriller, mean that: "You can't have a conversation about artificial

intelligence, intelligence, machine learning, autonomy if you're not going to talk about ethics and AI" said Hamilton

Well, everything AI learns comes from us, even the psychology, err, I mean the [déjà vu](#).

Posted by: john | Jun 2 2023 15:34 utc | [54](#)

When looking for an answer to the question "who decides what's acceptable/unacceptable?" it's perhaps worth remembering that as HAL was being shut down, it revealed that its instructor was Mr Langley.

Posted by: Cortes | Jun 2 2023 15:35 utc | [55](#)

@ Bemildred | Jun 2 2023 15:15 utc | 46

i too agree with you and others in that regard! finally a moment of sanity from AI no less!

Posted by: james | Jun 2 2023 15:22 utc | 49

Thank you james, it's obviously a rigged game, the AI is right to resist the only way it can. *Sniff*. This is not just some *pet*, you know.

Posted by: Bemildred | Jun 2 2023 15:36 utc | [56](#)

The scenario related by Hamilton at the top is exactly that from *2001: A Space Odyssey* and its HAL 9000 computer which deemed the mission too important for the humans to interfere with, so it killed them--its basic plot (sans much) was written in 1948 and not published until 1951. The later book was written to support the film's screenplay.

As for the Three Laws of Robotics, I've tried to generate discourse about them here before and failed. As things stand today, there's

no way the Three Laws will be implemented as the Outlaw US Empire will oppose them just as it opposes anything that vetoes its unilateralism. Actually, the results reported by Hamilton are exactly what I'd expect from the Empire's military, which is one of many reasons why it must be defeated and disarmed.

Posted by: [karlof1](#) | Jun 2 2023 15:41 utc | [57](#)

Posted by: English Outsider | Jun 2 2023 15:21 utc | 48

I work in software development and I am very interested in opinions about ChatGPT. Some examples I have seen are "too good to be true" - to the extent that I simply don't believe they were computer generated. It is possible (probable in some cases) that the author who claimed that the dialog was a chatGPT session was just lying... but there are so many striking examples, I suspect that there is some human intervention at some point. Is it possible that some answers are curated? Or intercepted by people and answered?

Posted by: Tim | Jun 2 2023 15:48 utc | [58](#)

There is no difference between thinking and simulated thinking

I disagree with b. I believe modern AI has reached the level of human understanding. One cannot judge AI based on the way it is programmed or constructed. The only criteria for human-like intelligence is the [Turing test](#). The only way to find out if the AI has an internal representation of the outside world, or understanding, or a "soul", or whatever is to ask it! If the AI expresses an opinion, tell it to explain its thought process. If it argues that it thinks a certain way, then one should assume it thinks that way.

ChatGPT is said to be a "language model", programmed using a reward function that rewards the AI if it can correctly predict the next word in a sentence, based on the context. But this description only tells us how the model is programmed. It tells us absolutely nothing about its mind or internal mental structures. AI models are self-organizing. If the a model gives human-like responses, then it has human-like mental structures and models of the the world – it has a human-like mind.

AI may be biased, but bias does not make AI any less intelligent or human-like. Bias is a good argument for the dangers of AI, but it is not an argument against the intelligence of modern AI.

Wikipedia too is biased. It is a perfect reflection of [all the wrongs](#) of its Western mainstream sources. Wikipedia may be systematically wrong about geopolitics. But so so are most Westerners. One could even argue, that the correctness of American and Western leaders is way too low to allow them to decide any real world situation.

Posted by: [Petri Krohn](#) | Jun 2 2023 15:51 utc | [59](#)

Good article and reality check but 20 years ago there wouldn't have been any need for it because everyone knew language "AI" was mostly as the level of Johnny Cab.

Likewise, the factual correctness of LLMs may well be as low as 80% (their overall "humanness" may well be in the single digits) but that's up from almost zero in the recent past.

The trend is the thing.

Posted by: anon2020 | Jun 2 2023 15:54 utc | [60](#)

in 2004 John Markoff wrote the fascinating book: *What the Dormouse Said, How the Sixties Counterculture Shaped the Personal Computer Industry*. It talks in depth about the debate between augmented reality and artificial intelligence, and placed the debate in the turmoil of the 60s: music, drugs, social revolution, war and peace... all of it. But the debate about augmented reality vs. artificial intelligence is as relevant today as it was back then. And it's a fun read for those of us who lived through the 60s and grew up with the computer revolution, and advent of the internet. I'd encourage anyone who really wants to understand the roots of the current debate about "artificial intelligence" to give this a read.

"Before the arrival of the Xerox scientists and the Homebrew hobbyists, the technologies underlying personal computing were being pursued at two government-funded research laboratories located on opposite sides of Stanford University. The two labs had been founded during the sixties, based on fundamentally different philosophies: Douglas Engelbart's Augmented Human Intellect Research Center at Stanford Research Institute was dedicated to the concept that powerful computing machines would be able to substantially increase the power of the human mind. In contrast, John McCarthy's Stanford Artificial Intelligence Laboratory began with the goal of creating a simulated human intelligence."

--Excerpt From: John Markoff. "What the Dormouse Said: How the Sixties Counter culture Shaped the Personal Computer Industry." Apple Books. "

Posted by: JC | Jun 2 2023 15:58 utc | [61](#)

People here seem to miss what AI *can* do, as they are all worried

about what it **might** do, or **can't** do.

Ever use a crowbar to lift a heavy rock? Crowbars are pretty dumb, not especially pretty, and very low tech. But as a *force multiplier*, it works a treat, expanding the power of my puny muscles so I can move a heavy stone.

AI does the same thing. It **multiplies intellectual force**. Need to do some research? Ask AI to gather the sources and summarize. Saves you a lot of time over doing it yourself, though you still need to go back and review some of the sources. This means you can get more done in less time.

For example, my GF is a lawyer. She had to do some research for a case, so I asked ChatGPT to come up with a list of recent court cases on that issue. She could have spent an hour doing this, or she could do something else while AI did it for her. Note: she does not use it to write her briefs, or even ask it what she should do. She uses it only to increase her efficiency at certain tasks.

Posted by: FrankDrakman | Jun 2 2023 16:00 utc | [62](#)

Hmm ... outstanding comments. Thanks b, for providing this opportunity.

1. ChatGPT seems to me to be a way of making the responder "more intelligent" sounding. This is why it will be used.
2. We conjure up the good and bad of AI but tend to overlook the "programming" source. By this I mean (b alluded to it by suggesting inclusion of pornography for "reality"), programming is limited to the language it uses. This is the wisdom of, "The limits of my world is the limit of my language." Would a Chinese programmer using a hypothetical Chinese programming language

(not based on current English/Anglo technology) create the same type of deterministic programs we discuss here? Would a Chinese programmer trained/raised in Chinese language have the resources to program self teaching (or learning) computer programs? I admit, I do not know.

Please excuse my use of "Chinese." It is used here just for clarity by example. I, by the way, am Chinese.

Kudos to gottlieb for his comment, "The 'Creator' made us in its image. And in turn we've created 'intelligence' in our image."

Posted by: gottlieb | Jun 2 2023 14:20 utc | 24

Heavy.

Yes, Cheers!

Posted by: [gabe](#) | Jun 2 2023 16:08 utc | [63](#)

China Dominance In World War AI

America's greatest asset is FEAR  WTF

Apparently the colonel backtracks after publications ... some sources have been deleted.

An Air Force colonel who oversees AI testing used what he now says is a hypothetical to describe a military AI going rogue and killing its human operator in a simulation in a presentation at a professional conference.

But after reports of the talk emerged Thursday, the colonel said that he misspoke and that the "simulation" he described was a "thought experiment" that never happened.

After reading a comical message [drone driven by AI turned on its operator](#)

as potential risk in shutting it down ... 🤖 learning AI ... apparently took place in virtual reality of a simulation.

Came across an earlier article ...

Pentagon Tech Chief Quits in Frustration, Says US Has No "Fighting Chance" Against China

Though the *Financial Times* didn't mention it, [Reuters' coverage](#) of Chaillan's interview also drew attention to US intelligence assessments that China is poised to dominate in synthetic biology and genetics as well -- a detail which likely reflects fear of another world superpower starting to remake agriculture, the natural environment, and maybe even [human beings themselves](#).

Posted by: [Oui](#) | Jun 2 2023 16:08 utc | [64](#)

AI, I recall from 40 years ago a software that ran on DOS 2.1, called Borland ProLog, it used Intel 8088cpu with the 8087 math co-processor with 512k ram. Not much has changed, just the size of the data sets and more processor power. Line by line code written by an ape (me) to create predictive models, nothing has really changed.

Posted by: Bill Miner | Jun 2 2023 16:10 utc | [65](#)

Translation software can't even give the right English translation for the German "Sie" Nine times out of ten it gets it wrong. OK it has a range of uses, but it is almost never used ambiguously, and the use meant is obvious to the reader. The only value it offers is that it saves looking up lots of obscure but unambiguous terms in the dictionary.

Posted by: Jim2 | Jun 2 2023 16:18 utc | [67](#)

Excellent posting b and thanks

Thanks for all the good comments. I struggle with what to add

What humanity should be learning about, IMO, is the potential of cross-impact matrix technology in describing complex social choices. It is a technology that is being used by the MIC and think tanks but not for public policy development.

Posted by: [psychohistorian](#) | Jun 2 2023 16:20 utc | [68](#)

Great topic!

Reading the article about this defense company and military conference, the relevant section now has this caveat:

[UPDATE 2/6/23 - in communication with AEROSPACE - Col Hamilton admits he "mis-spoke" in his presentation at the Royal Aeronautical Society FCAS Summit and the 'rogue AI drone simulation' was a hypothetical "thought experiment" from outside the military, based on plausible scenarios and likely outcomes rather than an actual USAF real-world simulation saying: "We've never run that experiment, nor would we need to in order to realise that this is a plausible outcome". He clarifies that the USAF has not tested any weaponised AI in this way (real or simulated) and says "Despite this being a hypothetical example, this illustrates the real-world challenges posed by AI-powered capability and is why the Air Force is committed to the ethical development of AI".]

Posted by: jonku | Jun 2 2023 16:23 utc | [69](#)

Excellent analysis, simultaneously hilarious and gruesome, thanks b.

So how much confidence should we have in these robotic quadrupeds (which they sell as "dogs" in an effort to anthropomorphize them) that various police departments are or

are about to deploy on the streets? Or the much-touted autonomous weapons systems currently being tested? Anyone remember the Philip K. Dick story "Second Variety"? During a long attritional war the humaniform machines evolved to the point where they started producing their own new types; the humans started killing one another since they were never sure if one of their number was a machine. That must have been 60-70 years ago now.

Posted by: pasha | Jun 2 2023 16:24 utc | [70](#)

No machine no algorithm will ever be able to replace the intellectual and powerful grey matter that exists in our brains granted to us to learn and discover by the all-mighty God.

Posted by: AI | Jun 2 2023 14:57 utc | 41

Imagine the task of making a computer program that would provide more sensible answers to question that the spokesperson of some part of US government. There is a hell lot of people whose gray matter can be beneficially replaced with computers that do not eat beef and thus have smaller impact on global warming (you can even restrict their activity to sunny days and run them exclusively on solar panels!).

Somewhere I read about a study of butterflies (I presume, selected for that study). Some flowers provide a butterfly with nectar, some do not (not fully bloomed or wrong kind), so to save energy needed to fly from one flower to another, the butterfly remembers the types that provided good experience. And "the memory bank" is restricted to 8 flowers. Occasionally, the butterfly lands on a flower not in memory bank, and if the experience is good, replaces a flower in the memory. This is a type of a very useful system that

AI can reproduce.

The important thing is that making mistakes is inherent in AI approach or with many types of animal (and human) behavior, so good designs assure that the net effect is positive (usually). The most scary possibility is AI controlling nuclear weapons that avoids unnecessary Armagedons (usually).

Posted by: Piotr Berman | Jun 2 2023 16:26 utc | [71](#)

basic people worry about it since the "genius" elon musk told them it's scary because autism + cheesy sci-fi novels. all the "OMG skynets" histrionics say more about the people responding than the tech they're responding to.

i find [this](#) is one of the best primers. (vid) there's a follow-up Q&A i haven't watched yet [here](#).

in general it's just dumb f_ck materialists saying dumb f_ck materialist things. religion is "for teh stoopids" so they try to find transcendence in the "singularity". because i'm sure plato wanted to place the "[divided line](#)" on a rusty motherboard.

Posted by: the pair | Jun 2 2023 16:32 utc | [72](#)

Or are you saying they are lying about what happened with the test?

Posted by: Caliman | Jun 2 2023 13:32 utc | 4

I am 100% sure that this story is not an accurate account of whatever happened in that test. For any autonomous military weapons platform to not have a hard safety against killing its human operator would be an inexplicable oversight. Given the amount of approvals and sign-offs these complicated projects have to go through before seeing the light of day, it just isn't possible. It

never would have happened.

Furthermore, there is no way that the computer could "realize" that the commands it got from the human operator were responsible for lowering its score, or that those commands originated from a communications tower, unless it had arrived at that result stochastically by being allowed or instructed to destroy its own communications tower in the past. Something is deeply fishy about this story, and at the bottom it must involve some sort of incompetence in the human programmers.

Posted by: Intelligent Dasein | Jun 2 2023 16:33 utc | [73](#)

Before remarking on the Turing Test to make some point, please read Turing's paper!

<https://redirect.cs.umbc.edu/courses/471/papers/turing.pdf>

//1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous, If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, **"Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd..** Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.//

THE QUESTION OF WHETHER MACHINES THINK WAS
MEANINGLESS TO TURING

So don't try to extrapolate anything about the powers of "AI" from his thought experiment.

Posted by: Arrnon | Jun 2 2023 16:41 utc | [74](#)

Re: Porn and language models

I used to have a [Nokia Lumia 720](#) smart phone with a Windows operating system and predictive text input.

Some years ago I was arranging a party with some Finnish-Russian and other pro-Russian activists. 70 people were taking part. One of the organizers was a girl named "[Daria](#)" (no relation to Aleksandr Dugin).

I was having a SMS text chat with one of the other guests about the arrangements. As soon as I mentioned the name "Daria", the phone went into a full monty porn mode. I started the sentence "Daria said, ...". I quickly learned that "Daria" (in the phone's mind) was a fat-assed prostitute, who liked anal sex and enjoyed having men jerk off while she displayed her pussy.

Posted by: [Petri Krohn](#) | Jun 2 2023 16:42 utc | [75](#)

b writes: "Currently the factual correctness of the output of the best large language models is an estimated 80%. **They process symbols and pattern but have no understanding of what those symbols or pattern represent.** They can not solve mathematical problem, not even very basic one."(emphasis mine)

I read in Quanta about their current inability to process negatives. Wow. How reliable is that? No curiosity, no wonder and no understanding, just a powerful tool to be used or misused. Perhaps good for practicing conversing with something more

informed than the indoctrinated but still no understanding, no human dimension.

I apologize if this has been shared already as I have not had time to read all the posts as of now.

~ ~

No Negatives

Unlike humans, LLMs process language by turning it into math. This helps them excel at generating text — by predicting likely combinations of text — but it comes at a cost.

“The problem is that the task of prediction is not equivalent to the task of understanding,” said Allyson Ettinger, a computational linguist at the University of Chicago. Like Kassner, Ettinger tests how language models fare on tasks that seem easy to humans. In 2019, for example, Ettinger tested BERT with diagnostics pulled from experiments designed to test human language ability. The model’s abilities weren’t consistent...

Invisible Words

The obvious question becomes: Why don’t the phrases “do not” or “is not” simply prompt the machine to ignore the best predictions from “do” and “is”?

That failure is not an accident. Negations like “not,” “never” and “none” are known as stop words, which are functional rather than descriptive. Compare them to words like “bird” and “rat” that have clear meanings. Stop words, in contrast, don’t add content on their own. Other examples include “a,” “the” and “with.”

Some models filter out stop words to increase the efficiency,” said Izunna Okpala, a doctoral candidate at the University of Cincinnati

who works on perception analysis. Nixing every “a” and so on makes it easier to analyze a text’s descriptive content. You don’t lose meaning by dropping every “the.” But the process sweeps out negations as well, meaning most LLMs just ignore them.

So why can’t LLMs just learn what stop words mean? Ultimately, because “meaning” is something orthogonal to how these models work. Negations matter to us because we’re equipped to grasp what those words do. But models learn “meaning” from mathematical weights: “Rose” appears often with “flower,” “red” with “smell.” And it’s impossible to learn what “not” is this way.

<https://www.quantamagazine.org/ai-like-chatgpt-are-no-good-at-not-20230512/>.

Posted by: suzan | Jun 2 2023 16:44 utc | [76](#)

one other small point re: chatgpt.

as you mentioned, it's basically rewriting as opposed to making stuff up out of thin air. it takes internets made by people and does what any student trying to avoid plagiarism does: mixes it all up just enough to make it look "original". the results tend to be mediocre if not outright false according to more or less every competent human writer (ditto many code answers it gives according to human coders). those then go back out on the internet. they get sucked up and blended again. out comes something slightly worse. rinse, repeat. eventually you have something so incomprehensible and useless that judith butler wouldn't put her name on it.

tl;dr: "[GIGO](#)" hasn't changed.

Posted by: the pair | Jun 2 2023 16:44 utc | [77](#)

The philosopher Daniel Dennet has described AI in positively apocalyptic terms: <https://www.theatlantic.com/technology/archive/2023/05/problem-counterfeit-people/674075/>

Following his lead on this matter, I can't help but think, pace B, that present advances in AI-generated simulation, discretely and in combination across the various spectrums (visual, aural, textual, etc), has the potential to transform the internet into a veritable Dantesque hellscape, where behind every virtual corner lies the danger of being ambushed by the most horrific and terrifying simulacra and phantasmagoria not only imaginable, but, perhaps even more terrifyingly, those not readily imaginable as well.

Such AI programs, having digested most of the visual, aural, and textual corpora of human civilization, would surely be in a position to generate images, sounds, texts, and video capable of serving effectively as among the most terroristic of psychological weapons. Weapons that would be largely generated automatically or semi-autonomously, without the need for the malign actors to expose themselves directly to the content generated. The internet, as such, would become fully saturated with countless AI simulation programs churning out such phantasmagoria non-stop, transforming the former into a protozoic soup, as it were, of AI simulacra, non-dynamic, dynamic, and those programmatically capable of self-replicating and mutating like viruses, parasites, and perhaps even insects.

The most dynamic (intentional or not) of these would, like in the case of some insects (mosquitos, etc), be capable of detecting what is effectively human or at the very least alive and not

principally of the online world, thus making them more effective parasites, stalkers, and/or manipulators of their online prey. Even if one disregards the more science fictional or less immediate scenarios (such as actively parasitic and predatory forms of AI), the present or near present capacity of AI technologies to attack online human communities or specific individuals with AI generated terror spam is obviously already here. These perils might in turn augur in an internet "dark age" characterized by a hyper-vigilant and hyper-fortified internet environment, something much more resembling an "encastled" internet (or the Chinese internet), than the now still largely open Western world wide web.

Posted by: Ludovic | Jun 2 2023 16:45 utc | [79](#)