

Beethoven, McCartney and Grapes ©

Beethoven, McCartney and Grapes ©

By James Dean VanBuskirk

Copyright © 2026 by James Dean VanBuskirk

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise) without the prior written permission of the publisher, except for the use of brief quotations in a book review.

Published by James Dean VanBuskirk

First Edition: August 2026

Disclaimer:

The information contained in this book is for educational and informational purposes only. While every effort has been made to ensure the accuracy of the information, the author and publisher assume no responsibility for errors or omissions. Use this information at your own risk.

Author can be contacted at: jimbo202679@yahoo.com

To Begin

AI is poised to be the next intellectual revolution and it could change everything that follows it.

If this is true, I would like to know how that would occur and the timeframe within which it could happen. I would also like to understand the various mechanisms involved and how they might work to usher in this great evolution.

For this discussion I have chosen “music” as my entry point and I will explain why in Book 1, and then we'll progress to the more general topic of AI as it expands into all aspects of life in Book 2.

I have taken a lot of diversions along the way as you will see, each one hopefully, adding some understanding to the various elements contained herein.

If AI truly is a game changer for humanity's future, then it would necessarily involve discussing a huge variety of topics - far too large for any one book, so I have arbitrarily limited the scope and tried to simplify where possible. There will also necessarily, be some topics either not mentioned or only partially discussed.

You may agree with some of this content or you may disagree - and with this in mind, I invite any discussion you may want to have on any additions or subtractions you think worthwhile. My contact email is on the title page.

I have used the abbreviation “AI” throughout both for current assessments but also far into the future. Many have other progressive abbreviations such as “AGI” to better represent the future abilities which will undoubtedly be expanding but I kept the “AI” terminology throughout just for simplicity.

(AI = Artificial Intelligence / AGI = Artificial General Intelligence)

Contents

Book 1 – Music

Intro	6
What can a Computer do – and what can it Not do	6
What is Music ?	7
How to make Music – from Beethoven to McCartney	8
Magic	9
Current State and some History	10
Analog Recording	11
Enter the Digital Revolution	14
Magic Again – this time with Data	17
Computers by the Numbers	18
A couple of other Ideas	20
Grapes – Finally	22
Very Important Analogy	23
Computers and Music – Now	25
Magic Once More	26
Bubble ?	28

Book 2 - AI

Intro	29
The Great Grape Example	29
Upward and Wider	32
Whose Show is it ?	32
Duopoly	34
Some Computer History	35
Meet Andy	35
A Science Diversion	37
Andy again	41
Another Thought Experiment Please	42
Back into the Mud	43
DNA and More Complexity	46
Even More Complexity	47
The Big Jump	49
Where are we now and Where are we going	50
Conclusions	55
More Guesses	57
Some Dilemmas	57
More on Exponential Growth	58
Some Big Numbers	59
Coarse-graining	60
Data – How much	61
A quick Re-hash	63
Synopsis	71
AI Going Forward	73
Current	74
Short Term	75
Mid Term	77
Long Term	78
Very Long Term	80
A Personal Afterward	83

Book 1 – Music

Intro

With AI entering each of our daily lives and with new and sometimes shocking events occurring regularly surrounding this technology it is interesting and probably important to ask the question “where and how far can it go?”

This topic of intelligence, whether human or artificial, is not only difficult to define - it contains as many different facets of our existence as can be imagined. In fact, any “imagination” that can be framed as a thought is itself part of our intelligence and all this together is still only a minute subset of what we mean by intelligence. What the hell is intelligence?

In the attempt to answer this mammoth question I suggest it may be beneficial to focus on one category initially rather than on the field as a whole, trusting that if we can make some solid discoveries in that field it may generalize enough to aid our overall understanding.

To be more specific - I propose to first take a deep dive into trying to answer the question - “can a computer create music that is as good as real humans?” Or more generally, “can a computer truly be “artistic” given the full rich meaning of that word?”

We have hopes and dreams. We experience joy and sadness. We laugh and we cry. We experience thoughts and feelings on such a grand scale that language does not even have words to express all of these - so we write poems and music to expand our expressions beyond the dictionary's limits. Music contains and can elicit all these artifacts of human experience and more. It is intimately personal and yet experienced by all. Music is universal, it has been spontaneously created and consumed by all cultures and has evolved along with the technology used to create and consume it. It has been a force in human life throughout our entire history and is part of everyone's life in some unique way.

The digital revolution which we are in the midst of right now is changing our relationship with music in entirely new ways not only in its avenues of consumption but perhaps more importantly in its method of creation. This makes music the perfect entrance point for our quest to understand the more general topic of AI. This, of course necessarily involves computers and how they could possibly become intelligent. This appears to be an impossible task since we can't even define intelligence, but let's give it a try.

What can a Computer do? And what can it Not do?

A computer is not a person. It is a pile of silicon plugged into the wall. It does not breathe or enjoy the taste of a grape. If it ever does write an original lovely song I want to know how it did that.

So far in the entire history of this planet the only entity that could do that task was a human.

Is there a way to understand and explain a purely computational procedure that can somehow appear to match human creativity? It's not enough that a computer can have a larger memory than a human can or that it can access it faster - these are impressive facts but are not artistic or creative.

Deciding if a computer can bridge the gap into “art” will be our goal, and if it can, exactly what does it do to get there? I'm using music as my focus here but any field that can display artistry could be used. I will leave it to you to generalize to another field if you so choose.

Since we are perusing the musical abilities of AI, both current and future, that must include some understanding of computer technology.

Music has always evolved just due to normal unpredictable physics of humans - and computers have also evolved. The speed of computer technological progress is however exponentially faster than human evolution and we must include this phenomena and try to understand its effects. Music creation, production, and distribution is now evolving at computer technologies pace and this mismatch of interface tempo will be critical to our understanding. These are important and formative components in the question of whether AI is up to the task of matching human quality in the creation of music and we must include these concepts in our study.

Digital technology's penetration into all aspects of music will be better illustrated by a more complete knowledge of what came before. By understanding pre-computer technology, including analog sound recording, it will be easier to visualize the massive changes computers have energized and empowered.

For our purposes I would like to use Beethoven and McCartney as examples of famous musical talents, one active way before computers and one active during and after their presence in all aspects of music. These are just easy placeholders - you can sub in your favorites in both eras.

What is Music?

To assess this question deeply, we should probably begin by defining what we mean by “music”. In answering this question, it will likely be important to understand how music is made, how it is accessed, how it is judged, and even how it saved for future listening. These questions will necessarily also involve the associated technologies in each of these categories especially for the current question of whether AI and its computer parents can create music or only copy it.

Since we are asking if computers will be up to this task, we will also need to do a bit of a deep dive into current computer abilities and technologies. Let's go.

Obviously music is sound we can hear. Somehow it enters the air and is intercepted by our ears. We can probably agree that the ear-brain interface is also necessary since we are aware that we are hearing something. What we hear as music could also be considered “data” in the sonic realm and therefore producible and reproducible in many ways. In order to be called “music” it must contain some pitches and timbers presented in some sort of timing sequence - this is just to differentiate music from speech or just random sounds. We currently assume some human creator is organizing these

available sounds in some order with their creativity.

The range of sound choices used to be instruments or voices in Beethoven's day but now we have added to that palette basically anything that can make a tone - this essentially adds the computer to the mix of instruments. It will be important to get a handle on the range of elements and perfection the computer has added to better prognosticate where developments can take us in answering our original question.

If we are asking if computers (AI) can create real human quality music, we had better know what technology can currently accomplish in order to better guess what is next in the timeline.

How to make Music - from Beethoven to McCartney

Easy! We can play it or sing it. We can also write it so someone else can play it in the future, or these days we can also record it although this was not available to Beethoven. He could play any one instrument at a time but he could only write whole symphonies in his head. As musicians we learn to pre-hear some sounds and pitches after becoming familiar with our instruments or do ear training to further advance this inner awareness.

Remember this pre-hearing concept as it will become very important later. As a musician becomes more experienced in the playing of their instrument they can hear what a certain finger position will sound like before they play it. A trombone player must learn where to adjust the slide along its travel to produce the note they want. At first it's experimental but by repetition it becomes a learned behavior and becomes the bedrock of being able to play that instrument.

The most obvious example of pre-hearing is when we sing, you try to hit a particular note and adjust what you hear with what you want to hear in real time. Some can do that adjustment quickly and remember how to do it for the next note - we call them "good singers." We learned to talk this way too. As babies we adjusted our tongues and vocal cord muscles so that mommy would understand when we said "cookie." It's actually just a feedback loop. Create some output, watch to see if that output accomplished your goal, and adjust that output until it does.

Beethoven had advanced in these skills to an incredible level. Since he was almost totally deaf by age 26 and continued to write symphonies for the rest of his life this proves that he could pre-hear these entire symphonies of instruments all in his head. And for Beethoven, while he could still hear, he could only hear these ideas in the real world if an actual symphony would really play it.

Fast forward to today and a computer can read that Beethoven score and with a simple MIDI (Musical Instrument Digital Interface) connection, duplicate every instrument sound and playback that score without error on demand. The ability to record music was the first tech innovation to impact music in a huge way. Now you could hear music without the need for a real person to play it and player and listener could be separated in time and space. No longer did you have to assemble 60 people together with all their instruments at the Symphony Hall at 7:00 p.m. on Saturday to hear some music. At that time if you wanted to hear that symphony you had to be there too.

This making of music involves not only the production of some sounds but also the sharing of those sounds either live for Beethoven or now live or recorded for both Beethoven and McCartney. This sharing doesn't yet speak to popularity or enjoyment but we will go there soon. If we're going to be asking if the computer can make music as good as humans we will need some way to measure how appreciation can be judged. This will involve defining the range of human cognition and comparing that with the ability of a computer to create sounds we interpret as "music". Further, we will need a way to discuss that tricky notion of "appreciation" more deeply.

Magic

To do this I want to introduce the notion of "magic". There is "something" that is "somehow" contained in Beethoven's music and also in McCartney's that is demonstrated by the easily observed popularity and longevity of that music. It is probably impossible to define what that quality is but the music is clearly appreciated by large numbers of humans. They buy it, memorize it, and want to hear it again. They may smile or hum along, they may just tap their foot but they clearly and observably demonstrate that they like it.

Whatever that element is that causes large numbers of people to enjoy it - I want to call that element "magic". Instead of getting hung up in trying to define what is likely not definable, I'm proposing that by looking at its effects on people who are listening we will eliminate the need to define it in words, at least for our current study.

We now have good ways to judge whether a piece of music has "magic". For Beethoven it was the number of orchestras that were willing to play his pieces and how the audience reacted. For McCartney you can monitor sales of concert tickets and recordings and even more immediately count the streaming numbers from all the current services. Nobody in either era was required to desire this music, they were just somehow "magically" attracted to it in great numbers - giving us a metric to use in judging the appreciation level.

So how do Beethoven and McCartney make this music that we think contains magic? I can speak more knowledgeably on the McCartney era than the Beethoven era but I'm guessing most elements in this diagnosis will apply to both.

So McCartney grew up with musical parents and had occasion to listen to the music in his environment - some recordings, some radio, some concerts, and apparently some live rehearsals in his living room of his father's bands. He had a piano in the house and a guitar procured early in his life. Somehow he apparently liked music since it is obvious he chased the dream.

His "magic" was his ability to recombine and synthesize all his listening experiences and any techniques he could apply from his available instruments (including voice) and either copy what he had heard or recombine elements into something new - both of which he did in his bands. If it was pleasurable to him he may just remember it or write it down and soon in his developmental years record it to better and more perfectly recall the components of a soon-to-be publicly available song.

This creation process could be done mentally and even pre-heard in his head, but until it is shared and available to others we can't make any judgments. How many symphonies were written in

minds only and never denoted or performed? An interesting philosophical question but too off topic for this treatise. For right now therefore, let's agree that music has to be performed or somehow shared with our ears in order for us to decide if we like it and whether it contains that “magic” thing.

We now have at least a little handle on what “magic” is and how we might measure it. Since our goal here is to decide if a computer can reach that lofty goal, we should understand what can be done currently with computers and how they got there.

Current State and some History

Let's do a little survey of what computers can do right now as well as some history related to their application to music. First I would like to explain four current technologies that will create a framework to understand more fully what computers are already doing in music and their impact on music production.

1. Drum machines.
2. Sampler technology.
3. Digital Recording vs. Analog
4. Suno - Text to Music

----- Number one is the drum machine. Drum machines were the earliest of computer music devices to become commercially available. Apparently the idea started way back in 1931 when Leon Theremin created something called the “Rhythmicon”, Of course it couldn't become useful yet because computer technology was non-existent at that time. Fast forward to the 70's and 80's and I actually purchased one of the earliest drum machines (the e-Mu). Since drums are a limited palette of sounds that need to be played in some repeatable and measurable timing sequence it was the easiest entry point to initiate the computer into the music game.

The e-Mu contained digital sound files that would mimic the sound of each drum and record and playback those sounds with reference to a time clock. It could be adjusted both while recording and during playback. This was an incredible new method of mimicking what a real drummer does. You could set the time clock and just push a button when you wanted to hear a bass drum in that time sequence. Once you had that recorded in a way you liked, you could add a snare drum over that bass drum part until you liked the combination. Add other drums one by one including cymbals and hi-hats...and you now have a full drum kit without needing those pesky drummers.

You could add any drums at any time or remove any part or even a single note. You could record at a slow speed a part that was difficult or impossible to play at the speed you wanted and just turn up the tempo knob. Even on the earliest machines you could also just add a drum sound one by one wherever in the measure you wanted that sound to play and build up complex parts one piece at a time without ever having to actually “play” it.

These early drum machines sounded sterile since the individual drum sounds were just digital files that someone created that sort of sounded like a particular drum and they were also played with perfect timing which of course no human can match. Nonetheless they were amazing to experiment with and basically were the catalyst that launched rap and hip-hop music into public awareness.

Soon, of course this technology from the early drum machines exploded into the whole field of recording allowing exponentially more options to play and record music.

Each element of that early drum machine was then improved upon. Each drum in that early machine would lie perfectly timed in your recording. If you were playing 16th notes each note would be placed perfectly in those time slots even if you couldn't play them perfectly. This was called “quantization” and initially sounded fantastic. Soon however it was obvious that it didn't sound real and was too perfect. Also the individual sound of each drum sounded a bit fake. Enter the number two technology we should discuss.

----- The second advancement we need to discuss is sampler technology. Let's examine how a recording engineer would attempt to record a real human playing drums. First you would get the best drummer and the best quality drums available. Next you would put those drums in the best sounding room or area that you could find or create and then you would find the best sounding microphone. You would then record that drum and experiment with placing that microphone so it would yield the best sounding “whack” that you would then capture on your recording tape.

Since the best tape recorders of that day usually recorded at a tape speed of 15 inches per second our two-second recorded bass drum sound would occupy 30 inches of tape. Now let's rename this recorded bass drum sound a “sample” and realize it is actually just data on a tape representing our best try at recording the best sounding real bass drum we could accomplish.

Since we now have a computer, let's transfer that sound “sample” from that 30 inches of tape onto a computer chip and put it in the drum machine and have the drum-machine pad trigger the playback of that sound. Now the sterile digital pretend bass drum sound we originally had is now a real drum sound recorded with the best equipment and care available at the time. It would be impossible to trigger 30 inches of tape but easy to trigger that chip containing the sample. Computer to the rescue - adding an entire new way to play drums that sound real because they ARE real.

More to come on this sampler technology but for now let's move onto the third technological development we should understand.

----- The transition from analog to digital technology has impacted everything in general but the entire musical field probably as much or more than any other category. In order to illustrate the changes that recording technology brought to the music industry and music production we should first understand how analog recorders work.

Analog Recording

Analog recording is a method of storing an audio signal directly onto a physical medium by creating a continuous proportional replica of the actual sound wave. The first analog recordings were actually done on a wax record with a cutting head cutting grooves that match the sound waves being recorded. Recording using tape began in the 1930's in Germany by BASF and variations and improvements to that design are still used today although many have been replaced by digital recorders

in the 1990s.

When using an analog tape recorder, the original sound waves first encounter a microphone which has a metal diaphragm that vibrates sympathetically with those waves. The moving diaphragm is mounted in a magnetic field and creates an electrical signal that travels down a wire to a recording console. Here the signal level is adjusted and tonalities tweaked for the best sound. A compressor/limiter is often inserted in the circuit here to ensure that the sound level does not get too high and overdrive the electronics.

The electrical signal is then routed down another wire to the tape deck. The electrical signal is fed to a tape head which changes the electrical signal back into a magnetic field and rearranges the iron oxide coating on the passing tape to mirror the original sound that entered the microphone. Ultimately, that encoded sound makes it through an amplifier and then to your speaker which hopefully faithfully recreates the original sound waves which first encountered the microphone diaphragm at the beginning of this process.

An analog sound wave is transmitted through the air at about 700 miles per hour and is actually a series of compactions and rarifications of that air recreated by the movement of your speaker. When the speaker cone is moving towards you and away from you it is essentially creating a pressure wave that once again hopefully mirrors the exact wave that was originally recorded.

If you graph those waves as they travel through the air, you get the famous sine wave shape that we all love to see and hear. Make a mental picture of that wave for a moment, it will be important later.

The actual process involved in creating a record is different depending on what you are recording. Our example here will be a band, maybe even The Beatles recording a song.

One way to accomplish the act of recording is just play the song. One microphone can do this task and can be routed to one recording head on a tape deck. When played back, this recording is “mono” (one channel) and a few of the early Beatles recordings were actually done this way.

“Stereo” sound requires at least two tracks on the tape deck since the left and right channels supposedly more accurately mirror the slightly different signals that arrive at our two ears. Those two channels are also necessary for us to hear the left and right placement of instruments on the soundstage.

If this was your recording procedure there are still some harsh constraints that remain – for example, everyone had to play the song perfectly and all at the same time. The microphones you were using had to be optimally placed on each instrument and voice and had to be balanced perfectly. You could adjust all these individual parameters as you were recording but not after.

A new set of amazing options soon became available with “multi-track” tape machines that eliminated many of these constraints. By stacking the recording heads vertically you could record multiple tracks separately - not only fed by separate microphones but could also be recorded at different times. Stacking the recording heads was a technological challenge but was developed quickly and revolutionized the process of recording. Two, four, eight, sixteen then thirty-two track machines were produced over the ensuing years.

Ringo now can have five or six microphones placed ideally on each drum and cymbal. Each mic can have its own track so later on each can be individually adjusted for tones and volume. Paul has a track for his bass. John and George each have a track for their guitars, also all individually adjustable.

They could still just play the song together at the same time and attain far better sound quality due to the individual adjustability of the multiple tracks but they also had the option to play separately until each had his optimum performance. Paul can play bass on one track and add a guitar later. He could also re-record his original bass part if he has a better idea after hearing the song as it develops.

The multi-track tape containing all those optimal tracks was then fed back into the recording console where each individual track could be optimized for sound quality and “mixed down” to a two-track tape and ultimately pressed into a stereo vinyl record.

There were still some limitations however. Each track had to be recorded by real people in real time. Redos could be done of course and even some parts of each person's performance could be slightly repaired by “punching in” parts and at other times actually cutting the tape and inserting another section. Hmm...”cutting and pasting” - sounds familiar.

Another problem was that the tape could easily become over saturated with a signal that was too loud and distorted the sound. In order to get the best end result sonically, it however was important to have the instrument recorded as loud as possible to overpower the background noise. Compressor/limiters were developed that would monitor the incoming signal and reduce the level instantaneously if the signal was too loud.

Now you could record with the record level set quite high and just compress off the peaks. This was also used so radio stations could always be broadcasting a loud signal. If you were searching across the dial for a station to listen to it was more likely you would stop at a station you could hear. If a song starts with a quiet guitar intro, that intro needed to be full volume to catch you – then when the main song kicks in and all the instruments are playing full-volume the compressor reduces that volume back to the allowed broadcast levels. A great tool for recording and broadcast safety but the downside here is that the dynamic range of the music enhances its effectiveness and it was almost eliminated by compressing all the peaks. Most records and even early CDs use compressors when recording due to the oversaturation problem and pressure from radio broadcasters.

The multi-track tapes were also huge, heavy and expensive. The stacked recording heads got up to two inches tall requiring long reels of 2 inch wide tape and an expensive machine that could feed that heavy tape along at exactly 15 inches per second. The iron oxide coating on the tape also had to be perfectly manufactured and it would wear off if you passed it over the record head too many times.

Even with all these limitations you can imagine how these multi-track machines drastically altered the process of recording and allowed many more options to exist that had never been available previously. Now thousands of people could listen to that song for free on the radio or even have a record in their house to play anytime.

This recording process is a perfect example of a new technology that birthed an entire set of

artistic procedures leading to a business model that had never existed before. An incredibly transformative time when a set of relatively simple technological inventions change the entire field of music and foreshadows what is to come.

----- Now that we have a handle on the analog era let's see what effect digital technology had on recording procedure and further into the entire sonic realm.

Enter the Digital Revolution

First we must begin by understanding how digital encoding of a sound wave works. That 700 mile per hour sine wave we spoke of previously enters an analog-digital converter and in each microsecond a sample of that wave is taken in essentially the same way our drum machine sample was made. The difference here is, with the technology developed since my eMu drum machine, that sampling is now occurring at 44,100 times per second (for CD quality sound). These individual samples are then arranged in order and when played back can recreate to an amazing accuracy our original sine wave. It's not a perfect replica of course because its actually small packets of sonic information stacked side by side in the time domain, but with the sampling rate being over twice the frequency human ears can hear, it was assumed to be more than adequate for accurate reproduction.

This sounds like just a minor switch in the method of encoding and saving the sonic information of a song but let's look a little bit closer.

Remember that new digital drum machine chip with the bass drum encoded? When you trigger that chip we now call that an “event”. That event can be moved anywhere in the timeline by copying, cutting and pasting. Back in the analog days that was marginally possible but very imperfect and time-intensive - now it's as easy as in your word processor.

Placement of any event in the timeline can be chosen with near perfect precision and adjusted easily. In the analog realm you could speed up the tape but that also raised the pitch of what you recorded - not so in the digital realm. Speed was just how fast the events happened and did not affect pitch. Timing is just the placement of digital “events” and the sounds are just files (samples) that can be adjusted and triggered easily.

In the analog era we had that tape passing over the tape machine head to retrieve the magnetic information that was encoded in the iron oxide coating and that encoding was always imperfect. This imperfection resulted in tape hiss and was a major problem. Both DBX and Dolby Labs invented technology to reduce this problem that actually worked quite well but it introduced other artifacts that were more difficult to solve. Also the tape had to travel over the heads at a perfectly constant speed and when ultimately pressed onto a vinyl record, that disc had also to spin at a consistent speed. In addition, the needle on your record player passes over the vinyl grooves at different speeds in the outer grooves then the inner grooves and we all remember the sound of static and scratches from dirt and a worn needle.

There was also a problem with low frequency analog sounds when pressed into the grooves of a record. The lowest note on a grand piano is 27.5 cycles and McCartney's lowest bass note was 41.2

cycles. The range of frequencies for human voice is about 90 to 300 cycles and the cymbal sizzle frequency is in the thousands or tens of thousands of cycles. (Humans with good hearing can hear from 20 to 20 thousand cycles).

Remember, the speed of sound in dry air is about 700 miles per hour - if you do the math, McCartney's lowest bass note has a wavelength of 41 feet, and Ringo's cymbal wavelength at 10,000 cycles is 1.3 inches.

If you were trying to represent both those frequencies physically pressed into the groove of a vinyl record, the low frequency variations would be so large that they would cause the needle to literally pop out of the groove. Early records sound very thin because they had to keep the lower frequencies muted to allow the needle to track properly. When more aggressive music like rock proliferated and you wanted full powerful bass, the RIAA (Recording Industry Association of America) came up with a solution that is still used today.

They created a “curve” that became an industry standard whereby when pressing a record the lower frequencies were lowered and higher frequencies raised. This allowed the needle to stay in the groove and track all the frequencies more properly. When played back the amplifier had internal electronics that would reverse the curve and boost these lower frequencies back to the original level. This is why if you have a separate turntable it must be plugged into the special “phono” input so the RIAA curve can be re-implemented.

All these problems inherent in the analog realm disappeared in the digital realm

The RIAA curve is no longer necessary since low frequencies and high frequencies are just information samples. You still have to make sure your amplifier and speakers can handle these low notes but the jumping needle worries are gone.

No tape hiss and no spinning chunk of vinyl that has to revolve at the perfect speed - and no huge reels of tape that need to be perfectly fed across the recording head.

You still need to honor the fact that the stacked digital samples have also to be encoded and played back with impeccable timing - but with computer clocks that timing is far more accurate than spinning plastic. Digital penetration into music then spread quickly.

In the drum machine example there are currently many thousands of digital samples of every conceivable drum or percussion instrument recorded initially with all the care previously mentioned. A person or studio can buy a small rack mount device that contains thousands of samples to be triggered by anything that has switches. Many people have a digital keyboard where each key is a separate trigger that can recognize how hard and fast you push it (velocity). You can then easily map any of these drum samples to any keys you want using MIDI technology. Your keyboard is now a drum machine....

Generalizing past the realm of drum machines and taking advantage of your digital keyboard why not record a Steinway grand piano with the same care as we did with the bass drum. Use great microphones placed ideally around the room to pick up room ambience and record a sample of each Steinway note individually. Make each note into a separate sample and have each key on your

keyboard trigger a playback of the appropriate file.

Voila! You now have a pristine Steinway grand piano sound with your lowly \$600 digital keyboard.

Currently every possible instrument you can imagine has sample banks you can access – trumpets – violins - piccolos - Hammond organs and of course all the orchestral instruments. These available samples are an attempt to reproduce all the instrument sounds that have historically ever existed. To this we must add the non-instrument sounds that a computer can generate which is essentially almost any sound you can imagine. With modulators, multi-tonal mixing, adjustable envelopes, wave generation and more tricks any current synthesizer can and will surprise you with near infinite sound generating capacity. Listen to the sounds in the electronica field to get a taste of this.

Digital recording has revolutionized how recordings are done. Timing is now just speeding up or slowing down “events”. Pitches can be adjusted up or down without affecting timing. Individual “events”, which are actually sounds, can be cut and pasted almost as easily as you can accomplish that procedure in your word processor. Copying any part or combination of parts is easily done. With auto-tune technology you can even repair imperfect pitches in real time while someone is singing out of tune. Can you hear me Taylor?

Want a friend to hear your song or maybe record an added part? Just email the sound file to them and when they're done, they just email it back to you.

The manipulative ability that digital technology brought into the music world is profound, both in creating music and even more in recording it. My little 8-track analog studio that I built back in the day cost many thousands of dollars and suffered all the limitations of analog equipment. Today for well under \$ 500 you can have a DAW (Digital Audio Workstation) program on your computer that has unlimited tracks and plugins that can accomplish almost any digital sound enhancement you want.

You can record, mix and adjust every parameter, quantize or unquantize and edit note by note any part of any track. The DAW can also represent the actual sound waves visually on your computer screen making editing much easier. You can “see” when a note starts and ends making cuts and pastes much easier. You can grab any part of any wave and reshape it, retune it or move it. This feature adds your visualization skills to the procedure by literally “seeing” sound.

No tape hiss. No limitations on tempo or distortion - all recorded for free into your computer's memory. Send a file to YouTube or Pandora and now it is distributed worldwide.

Once again we have a new technology that deeply impacts the art of music in almost every way it is played, recorded and sold.

I think it's obvious that the computer's sound manipulation and recording abilities are already extensive and increasing its reach quickly - meaning it can copy, record, produce and reproduce sounds easily and with great detail.

Magic Again – this time with Data

Just great quality copying and reproducing however is not enough. We know both Beethoven and McCartney invented some music and liked it enough to share with other humans. Once other people heard it a great number of those other listeners also liked it and wanted to preserve it to hear again and even share with their friends .

These reactions are easily visible and easily measurable events. Of course buying a record is easy enough to measure too but let's go further.

There are certainly human reactions that indicate pleasure when hearing a musical piece - perhaps a smile or tapping your foot. Going deeper maybe some heart rate changes - deeper yet maybe even some definable changes in brain activity. Current science has mapped a lot of physiological changes in reaction to various events. FMRI machines can even tell which areas of the brain have increased activity with different stimuli.

This certainly is currently an imperfect realm of study, but who can argue that we don't know exponentially more about human reactions on more levels than we did even a very short time ago?

Relating this back to music, we arguably have increasingly accessible and identifiable data on how a person reacts to hearing music and whether they enjoy it. This data includes everything from smiling to toe tapping to buying the record and is all good evidence to believe a person likes the music. Add to this the quickly developing brain technology mentioned above and we can be even more sure of our diagnosis.

In previous years you would record your music and play it for your friends and then maybe take it to your local AM radio station to get some radio exposure. They would judge its worth by feedback from listeners which was slow and imperfect. Now, with the ease of recording and distribution you can get immediate and easily interpreted data on the popularity of your song. You know who is listening, when they are listening and which device they are using. In addition, the service doing the streaming knows a whole lot more about the individual doing the listening.

You have an incredibly complete data set that will show in great detail the listeners level of appreciation of any musical piece - and in real time.

This is essentially Google's business model applied to music. Google knows if you're looking for a light for your patio and sells that data to Home Depot. By monitoring how many people buy that light Google can decide how much to charge Home Depot for providing that new customer. This is just simple data management and is easily measurable with easy to collect data.

Basically by monitoring virtually everything about a person literally down to almost the atomic level - their physical appearance, voice, eye movements, excitement levels, brain activity etc. you can easily monitor a persons level of appreciation of a piece of music in real time and to an incredible level of perfection.

Access to all the data contemplated here is only partially available now but the storage, compute and algorithmic power necessary is being built daily in the AI data centers.

All the dystopian conjectures can come into play here. Cameras everywhere, phones in every pocket monitoring behavior, satellites always recording, smart meters on your house, computer chips in your TV, car and vacuum cleaner, a digital watch on your wrist. The staggering costs to build data centers and their associated computer processing abilities have no other conceivable reason to be created than to collect, store and process all this accumulated data. Much more on this later.

Back to the music field - this data will make it easier for the computer to decide if a person likes a certain piece of music. If you're unsure just collect more data! A video of your face would help. Monitoring your email and phone conversations to see if you have recommended that song to your friends would help too. It's likely much of this data is already being collected. Ask Google or the CIA.

It should be clear that this pile of data just mentioned cumulatively is very definitive of our ability to judge if something as simple as a song is being appreciated.

Let's now recall our earlier stated definition of “magic” - that unnamed quality that a piece of music contains that make people “like it” or “love it” and want to hear it again because its pleasurable. Different pieces of music for different people of course, but all must all elicit the above symptoms of enjoyment - and if it does, it therefore contains “magic”.

Conclusion - The computer can now decide if some music has “magic” in the same exact way we decide. It will look at the same data that we use; in fact it will have far more data than we could possibly have access to - and process it much faster as well.

Want to know if your friend likes a song? Ask him. See if they own the CD. See if it's on his Pandora playlist. Watch his face when you play it at your party. No problem! The computer has all this info and more.

We still don't know how to define “magic” but by monitoring its effects across all of humanity we can tell, to a great level of certainty if a song contains “magic” simply by monitoring the output of the listener.

This will become centrally important if we are trying to decide if a computer can be artistic and creative in the same way we humans are. So now we have AI that can determine if a musical piece has “magic” not by a real definition but by extensively monitoring our reactions.

Since we are conjecturing how this data and associated computer power will affect events in the future I think it would be helpful to put some numbers into the game. How fast can computers go, and how can they hope to process this mammoth amount of data they are collecting?

Computers by the Numbers

Basically current computers have two main sections: a CPU (Central Processing Unit), and a GPU (Graphics Processing Unit).

The CPU is fast. My flight sim computer is an i9-13900 which can essentially do over 200 billion operations per second. This is part of the computer that does what we normally consider as computing. It excels at math and serial threads of data handling. It executes programs and runs the operating system. Even though it has multiple cores and can do some parallel operations, it is optimized for doing simple complex tasks quickly.

The GPU is really the new kid on the block and where the magic lives. My flight sim computer has one GPU chip - the current best and fastest available: The NVIDIA RTX 5090. The GPU is purposely designed to handle the parallel processing needed in rendering complex visuals. It contains thousands of small specialized cores that operate on repeatedly identical tasks at the same time. The 5090 has 92 billion transistors and can process 3352 trillion AI ops per second !

I have three 4k video screens containing over 25 million pixels and each individual pixel can display 16.7 billion different colors and the screens can update continuously at 60 to 100 times per second. This massive data handling load is not possible with a CPU no matter how fast it could operate.

These new GPUs are the new component in the massive data handling requirements that any attempt at AI will need.

The 5090 chip that I have uses 95% of the electrical load in my computer and generates so much heat that I need five fans and a water-cooled radiator to dissipate the waste heat generated. Sam Altman (Open AI and Chat GPT) is planning on having 25,000 of these chips in his data center. Imagine all the water necessary to keep them cool and the massive electrical demand.

For AI these GPU chips are stacked into multi-level neural networks processing all the data we have been discussing and identifying patterns far beyond what a human could see. These GPU chips are being used in the reverse fashion – instead of creating massive amounts of pixels they are instead looking at massive amounts of data for any interesting emergent phenomena. New algorithms, look back options and Transformer technology (the “T” in Chat GPT) all add to the data handling and hopeful ability to operate with some kind of intelligence. The “P” in Chat GPT stands for Pre-trained and is essentially comprised of surveying and digesting all the collected data we have discussed herein.

Back into the music world - this pre-trained AI data has already logged every song ever recorded in its databanks and can monitor all new musical creations as they occur. Every song McCartney could have possibly heard and every response to that music he had, combined with every response that every person who ever heard his songs may soon be accessible data.

Of course AI data collection will go far beyond music responses but for our purposes here it's clear that “magic” is easily identifiable by looking at the massive datasets of collected responses. It is actually an easy computer task since we can look at every response from every listener. Feed all this data into the neural network, let the GPUs recognize patterns and you have the components of what works to create music we like. Remember, the “pre-trained” data is every sound ever recorded and every reaction to those sounds by every human who heard it.

Once the patterns are identified you simply shuffle them a bit and reassemble them into a song. Now you let some humans hear your new song and monitor all their responses with the data collection

procedures we just outlined. Take that data denoting all the responses and feed that back into the top of the neural network again. By doing that repetitively you have a foolproof procedure to continuously improve the responses to your song !

No need to know anything about music. No need understand how humans arrive at opinions. No need to have any notion of “magic”.

The massive data crunching ability of the networks herein described and the speed with which this is done combined with the massive amount of pre-training will create new assemblies of sounds that will be judged as “magical”. If it's not just feed it back through again which can be done thousands of times per second.

****Beethoven and McCartney both obviously had some human skill that created “magic” in music. A computer does not have that skill or anything like it; it's just a computer but by carefully monitoring every sound entering a person's ears, and also carefully monitoring any response of that person you can adjust the “input” to yield any desired “output”. Spooky but true.**

This is essentially market analysis by another name. We currently have test marketing specialists that test various products by altering the box it is packaged in or the taste of a cookie. They present the product to a group of people (the input) and measure the reactions of that test group (the output). The results they report will be used by the producer of the products to adjust the products and increase sales. Very slow and imperfect but still works as described. Slow because each test group takes time to set up and monitor and these groups are necessarily limited in size.

The computer is basically doing the same thing only with a world-sized group and perfect instant monitoring simply using continuously collected data. The processes just discussed and more are already in use now in all the AI programs generating language, video and music and answering questions posed to it that have never been asked before - by shuffling patterns as described.

As time goes on, the data amounts collected will continue to increase. Computer size and speed will continue to expand exponentially and the collection of data will reach further into all aspects of life on earth, including us humans.

All this collection of data is easily observed right now and will only increase the ability of the computer to mimic “magic” as measured by every observable event that can be used to identify it.

This discussion started with a focus on music but it has somewhat evolved into a broader discussion of AI, music just being a subset of that field. With a little generalization from this music focused study to the broader world at large it seems quite obvious that AI has some real tools to be creative and likely in some yet to be defined ways.

A couple other Ideas

With that in mind I would like to submit two other ideas to round out this discussion. The first is a little story from the past that gives us a possible hint of where things are going and perhaps how they will get there.

The first idea is actually a set of four events from recent history that will illustrate a few nuances on how the computer might expand on what it could accomplish in the areas we have been discussing. The second idea is best expressed as an analogy that I think may help our thinking about deciding on whether a computer can make new McCartney level music.

Let's start with the first of the four events to set the stage.

Event 1 / Alan Turing was a British mathematician who was an early experimenter with digital concepts and could foresee what might come of it. He worked during World War II on breaking the code of the “Enigma Machine” - the device the German army used to encode their secret communications.

About 1950 after the war, he was commissioned by the British Government to build the first digital computer which he called the ACE (Automatic Computing Engine). For this he is considered the father of the computer. It was his contention that the human brain is actually a computing machine and all brain processes could be done by a computer.

With his idea that computers might be able to exhibit intelligence at some point he came up with a test that would be used to decide if it had indeed reached that level. It was called "The Imitation Game" and involved a human having a discussion with two entities hidden behind a screen - one is a person and one is a computer. If the human having the discussion cannot identify which is which it was said that the computer passed the “Turing Test”

Fast forward to today and have a chat with ChatGPT and make your own judgment if computers have passed that test. It's pretty clear to me that it has easily passed that aging benchmark.

Event 2 / In 1997 Deep Blue, an IBM computer, handily beat the world champion Garry Kasparov at chess. Previously he had never lost a game. The computer had been trained by looking at thousands of games and could timely recall enough situations to beat the World champ. In chess there are 20 possible first moves, 400 possible interchanges over the second move and 10^{123} possible interchanges during a game. That number is ten with 123 zeros after it - a pretty sizable number and illustrates that with just memory alone computer technology could already better the best player in the world.

Event 3 / In 2016 Demis Hassabis wrote a program called AlphaGo which empowered his computer to play the ancient game of “Go”. Apparently Go is considered the most difficult game ever devised and is played on a 19 by 19 grid. The games were broadcast to a worldwide audience.

AlphaGo won four of five games. Hassabis had fed 150,000 games into memory and let AlphaGo play itself millions of times. AlphaGo was operating on a distributed network with 1920 CPU's and 280 GPU's.

After the wins he said: "Although we have programmed this machine, we have no idea what moves it will come up with. The moves are an emergent phenomena from its training." Could this emergent phenomenon be called intelligence yet?

Event 4 / This is the one that leads us over the edge and expands the range of speculation exponentially.

The next year, in 2017, Hassabis built another computer to play Go, and he called it “AlphaGo/Master”. He loaded into the computer only the rules of the game and the goal of winning but no actual games as training and let “Master” play itself over and over. At first the moves were often nonsensical but gradually it improved and ultimately beat AlphaGo.

This is revolutionary and opens a door that did not exist before. AlphaGo handily beats the world champ in 2016 with a computer trained on all previous games ever played. The next year 2017 Master beats AlphaGo 100 to 0 with no initial awareness of any previous games (no training data). Master was completely self-taught - this was in 2017 almost 10 years ago.

By training itself it could beat AlphaGo with moves no human had made before by recognizing patterns with neural networks and creating strategies that even AlphaGo could not recognize and that even 3000 years of human players had not yet discovered

This jump from pre-trained to no pre training is an important event in this timeline. A computer that beats the world champion, that was pre- trained on every game throughout history is an amazing event but completely understandable given that that is our method of learning,

A computer that somehow constructs its own methods of playing the game and then beats the one with all human history as it's database is stunning. Remember this progression from “AlphaGo” (pre-trained) to “AlphaGo/Master” (no pre-training). We will refer back to it later.

Is this emergent phenomena intelligence yet? This may be an important development relating to our discussion of AI in music due to the fact that the computer trained itself with no previous input of training data. This may also be a helpful skill in the quest to make unique music that exhibits our old friend “magic”

Computers are fast - and by feeding output back in as input we have a self-correcting mechanism operating at light speed. We do this all the time by watching reactions to what we do and adjusting our actions to achieve the desired result. We call this “learning” and it's our go-to process from birth onward. The computer can now do the exact same thing with wider scope and near complete monitoring with no sleep necessary.

As a point of interest, AlphaGo used 280 GPUs. Sam Altman (OpenAI), as stated previously, is now assembling 25,000 of them - and these are 2026 chips not 2017 chips.

Grapes - Finally

Just for fun let's follow the computer while it is collecting some data.

A young girl gets a grape from a bowl and places it in her mouth. She moves her jaw up and down a few times, smiles, reaches for another grape and says "mmm can I have more?" The computer can't possibly taste a grape but instead it encodes all these events into a searchable database. The

computer adds this to the same searchable worldwide database where it has filed away every human-grape interaction. From history alone the computer can now make quite certain judgments on whether that human enjoys grapes as it has recalled, almost instantly, every human reaction to placing grapes into mouths. It has also likely learned that moving your jaw up and down likely has a name: “chewing”, since most human movements seem to have descriptive names.

Let's go further.

The computer searches its worldwide human database of events involving the chewing of other items that produced the pleasurable results noted above. It then looks through its stored fact files and makes a list of chemical components in grapes. It also looks at all other pleasurable chewing events and makes a list of the components in those bites.

Time for the computer to employ its GPUs - so it then looks for patterns that may occur between chewing a grape and all other pleasurable chewing experiences from its vast database.

Aha! Sugar is a major pattern. Not the only one for sure but definitely high on the list. We already know sugar works but the computer knows exactly how well and for whom and far better than we could - and nearly instantly.

You can see where this goes. Employ feedback loops - each adds to the training data - experiment with as yet untasted combinations of chemical components - more data - self-correct routines that yield new combinations that are then used again as “input” and the process proceeds at light speed crunching all this new and speculative data. Maybe this yields a new grape hybrid or maybe a new round candy.... All this from a grape.

Very Important Analogy

I said previously that I would like to use an analogy that would be of some use in deciding about computers and music and indeed the whole AI arena. We have covered a lot in music production, history and technology in our quest to decide if AI can write a great song and this analogy may be very applicable in this scenario.

Let's say you have on a wall six feet in front of you an original Van Gogh painting. Next to that painting you have a video screen of the same size with a picture of that painting. It's probably easy to differentiate them but let's increase the pixels on that video screen. We can also adjust the colors pixel by pixel. Current 8-bit displays have 16.7 million colors; 10 bit displays have 1.07 billion colors; 12 bit displays have 16.7 billion colors available per pixel – that's a lot of color choices! Brightness can also be individually adjusted on a per pixel basis.

Well, “I can't see much difference between the Van Gogh and the screen copy from here about six feet away” but I can go closer and use my magnifying glass. Now I can see the difference - I can see the dots.

But of course I can now just increase the number of pixels until your magnifying glass doesn't work anymore!

This incredible degree of adjustability and the already available amount of compute power can, at some point, make that copy screen indistinguishable from the original analog painting. If you get a bigger magnifying glass I just add pixels. In real life you are not going to use any magnifying glasses and will do your comparisons from a few feet away or maybe even a few inches - but it's clear that the video screen copy will, at some point, be indistinguishable from the original.

So now the question is: Is there a difference once you can no longer tell? Technically, there certainly is a difference. The digital copy has a finite number of pixels and the original is a continuous medium, but again what if you can't tell? We may have to map the lumpy texture of the Van Gogh original to a not yet invented screen of some sort but its certainly conceivable that that could be designed.

We're not asking here can a computer do the original painting but the simpler question of whether the “effect on humanity is the same as derived from the copy”.

It seems clear that “functionally” there is no difference. Whatever meaning or appreciation you could get from the original, you can equally get that from that copy – given that is is indistinguishable from the original.

We have now unearthed an amazing concept - due simply to the incredible data manipulation abilities of a computer “finite data” equals “infinite data” at least as far as can be discerned by human senses.

This is one of those seminal events that we noticed earlier when multi-track tape recording entered the music field. A new technology that added completely new methods that had never before been available in human history and that literally changed everything overnight.

In our story we noticed this same thing occurred again to an even greater degree when digital technology entered the music scene.

I think the lesson that can be gleaned here is this: The “copy” of the Van Gogh only needs to be “good enough” that the observer cannot distinguish it from the original. We have now allowed “finite data copies” to stand in for our “infinite data originals” and be functionally equal in all discernible ways - We have also seen that a computer can do this.

That's a big deal especially for Van Gogh but even more for our quest to see if we can decide where computer driven AI can go.

I would like to name this phenomena that i have just described the “Van Gogh Effect” - any copy that is indistinguishable from the original is functionally equal to the original even though it contains far less data. We will come back to this idea later – many times.....

I need to go one more place before we try to summarize.

Computers and Music- Now

Since we have been focused on computers and music, we should survey where computers are now - mid-2026 . There currently are a number of text to music programs available that can literally do exactly that - you simply type in some “text” describing what music you want to hear. You can be very general and say, “something restful and sweet” or conversely you can be as specific as you wish. “I want a four piece band with a girl singer, three male harmony singers, i want two verses and one chorus and then a guitar solo all done in the style of the Supremes - key of A, tempo 120 and the words in this song should be about love lost”

Hit “Enter.” and in a few seconds you have two different versions of exactly what you asked for.

There are a variety of these services/programs currently available: Suno, Udio, Aiva, Beethoven, Mureka, and more. If you have not tried this go to suno.com and type some instructions. They have a free version and you will be completely amazed !

This is AI already becoming creative and starting another new epoch of everything musical. Single guitar players are using it to make backing tracks. Wannabe producers are experimenting with complete productions. It's happening now!

As the programs are maturing and finding ways to monetize their product you can now add functionality for a cost. You can separate the individual parts and recreate or edit them separately. If you like the whole song but not the bass, tell the computer to try it again. Adjust tempos, orchestration, lyrics or styles just like real humans would do - only its entirely a computer.

You can imagine all the legal ramifications bubbling up right now from this technology. The streaming platforms Spotify, Pandora, Tencent, YouTube, Apple, Amazon etc. are seeing about 40% of submissions that are AI or partly AI. It's still possible for these platforms to identify fully AI submissions but when it is a mix of AI and human its much more difficult. Some countries allow partly AI produced content to be copyrighted but it's still unclear on how to define it and even harder to decide who should be earning money for streaming plays. England is implementing a law - the “AI Act Alignment” - that would require a watermark on anything containing AI generated parts.

Even with this possible law it still begs the question of whether the producer will be honest about using AI. I think the Van Gogh analogy pertains here: what if nobody can tell if AI was used?

The quality of AI music generation is already undetectable to 97% of listeners and even the best computer algorithms written to detect its presence are failing in more and more cases. Middleman distributors like DistroKid and Routenote are trying to keep up with the technology but the current state of laws and requirements are a mess. The AI content in music is evolving way faster than our ability to manage it.

Again - this is yet another technological advancement that, in this musical case, completely changes everything.

Should the person typing in the text and getting a pleasant song as a result be considered the songwriter? How about being the producer? Should that person be able to receive royalties if they get

a million streams on Spotify? Even deeper: Should the writers of the copywritten music that AI used as “training” data be compensated? and if so how are you going to implement that?

With just a very little work you can now have the computer create an entire professionally produced song with lyrics and have an artist picture and bio in minutes all for free. A few of these deepfakes have already been caught by the streaming platforms after they have been streamed millions of times.

Where this goes is still undetermined even as the abilities of AI in the music field grow exponentially.

Magic Once More

Okay, we have now surveyed the musical arena looking at the long history of technological inventions and how they impacted all aspects of music: How it was made, how it was consumed, and how it is saved. We noted some of events and inventions that were turning points along that route all the way from Beethoven to Suno.

It's time now to answer our original question: “Can a computer create music that is as good as a real human?”

I think it is obvious that the simple answer is yes but with some important caveats.

Pop music is somewhat formulaic and tends to be rather simple in its structure and variety, the verses are basically repeats of the chords played and the timing. Since the computer has already digested every song available and certainly found often repeated patterns a recombination of these elements is quite easy. The incredible memory size now available and growing daily, together with huge advances on pattern recognition provided courtesy of the new GPU chips yields more range of available components for those songs than any human could possibly assimilate.

Just recombining these patterns is not yet artistic - but remember: The data collection of human reactions to each new song will easily show which ones create appreciation in the listener! Which “input” creates the desired “output” - that's all that has to be known. I think this process is already very successful as evidenced by Suno and the other current programs.

The data collection that is necessary to make this work is adequate today but increasing exponentially.

I do think there are some limitations however. If we move to more complex music like jazz for example, the difficulty presented to the computer is daunting. Much more complex harmonic structures, more time variations, more solos, and extemporaneous playing. Players inventing parts on-the-fly and only playing a certain part one time and never repeating it; even playing differently every time they replay a certain composition. Advanced jazz people can invent entire concerts as they play!

This is a lot more difficult for a computer to create than a pop song if for no other reason than there's far less training data that exists, however, I still think it is just a matter of more time and more

data collection. Millions of jazz performances have been recorded and can provide the training data and even though pattern recognition is far more difficult it will ultimately work just like our Van Gogh analogy.

Don't have enough data to currently create original jazz improvisations yet? Go get more data and monitor more output events until the listener can't tell the difference.

Once you have a handle on the tremendous amount of data collection coming and understand the complexity and speed that the new chips can handle - it's just a matter of time until real, new, never before heard combinations will appear.

Remember the AlphaGo to AlphaGo / Master transition we spoke of. AlphaGo had been pre-trained on thousands of games and beat the current champ. Master then was completely self taught with no training data at all and created moves that no human had ever made - and beat AlphaGo easily.

This phenomena will be coming to music as well and perhaps already is to some degree.

The Van Gogh analogy is forever at play here as well. The computer doesn't have to be perfect; it only needs to perform well enough so we can't tell the difference!

To understand this to a deeper level, we must not let ourselves be intimidated by the huge numbers generated both by computer speed and also by data handling capabilities. Billions of operations per second and trillions of bits of data handled and processed per second allows the computer to regenerate a variety of events as wide as the entirety of events humans have ever done but also in combinations not previously created.

By monitoring human reactions to these new combinations and using the output back-to-input loop to self correct, computer music will certainly be good enough so we cannot tell anymore.

We have been talking about various technological advancement over the course of music history that changed everything that came after them, including the ability of the computer to self-learn without need of pre-training on previous human compositions. This is a new door opening and may even “redefine” what we mean by music.

Where does this go from here? Here's a few observations.

The business end of music is already massively disrupted, that's for sure, and will entail lots of adjustments by all parties involved. The writing and production of music, is likewise disrupted even more at this point. We don't even know if a human made the music! This is an entirely new dilemma never before faced throughout the entirety of human history.

Music was always the exclusive province of humans. Not anymore! Changes cannot be any larger than that.

How do these new facts affect us individually? I suppose it depends on how much music is a part of your life. If you're not too involved and use music at most as background for doing laundry, the effect will be minor enough that you probably will not even notice the cataclysmic changes that

delivered that music to you. As we already learned, 97% of people already can't tell computer-generated music from the handiwork of humans.

Music is a much larger part of life for me than that. Having played it, recorded it and experienced it as a business I still derive great joy in the act of closely listening and appreciating its finer elements.

I have not yet heard evidence that the computer has reached the “pinnacle” of real creation of the more subtle elements of music, but I also can extrapolate that new abilities will come from the explosive speed of computer processing and the inclusion of GPU chips into the architecture. It's very close already.

As we have seen a “serial” computer alone could never be attempting the task of real musical artistry. The massive parallel processing provided by the GPU chips was the “game changer” in this arena. The main costs of data centers proliferating worldwide are driven by GPU prices and the attendant need for the power to drive them in addition to the water to cool them. These are certainly bottlenecks to development but are being managed one way or another. The amount of money, material and work involved is massive.

Bubble?

Many say this AI is only a “bubble” which it certainly is in many ways - but, you remember the story of the dot-com “bubble” from a few years ago? A lot of money was lost during that bubble too but time has shown that the underlying technology is now ubiquitous - everything is now dot com one way or another so the bubble event was just a scenario along the way - the underlying technology continued to flourish and has reshaped human life massively.

I think the current AI bubble is essentially the same. A lot of money will be made and lost but the underlying technology will remain and likely be the next phase change for humanity. Computers drove the internet and smartphones and almost everyone on earth was impacted in a very short span of time - and with GPU-enabled AI, the next revolution is already beginning.

Book 2 - AI

Intro

Let's turn to the big question: Will AI become as intelligent or more so than real humans? Our previous study of how it has impacted music I think, will provide some good relevant examples to help make this judgment. Let's do a quick review of musical history to highlight the changes and their effects.

Publicly available recording technology in the early 1900s allowed the saving and distribution of music to all. Multi-track recorders allowed one talented person to play many parts in each piece of music. Drum machines, which were the first digital machines, allowed anyone to be a drummer, even non-musician. As digital processors matured, sampler technology allowed any keyboard player to play flute and violin.

Digital recording was a huge leap and bankrupted most recording studios since anyone with a computer could now create and copy that music. Soon after that time with maturing deep digital capabilities involving GPUs, programs like Suno allowed anyone to create entire musical creations without any musical skill or knowledge.

So in less than 60 to 70 years music went from being a “human only” art form to a “human unnecessary” art!

It was important to spend some time on understanding exactly how the computer was able to create this incredible evolution in music - because future versions of AI will use many of these same processes.

The Great Grape Example

A lot of people when questioned about AI will say something like: "A computer cannot taste a grape and by missing that experience will never be able to be fully human." We have a hard time believing a computer can possibly process that because we have internal experiences and feelings and pleasure of eating a grape (or listening to a good song). I think this is correct but let's examine this claim a bit deeper.

It's certainly true a computer cannot taste a grape - it's a computer and has no mouth or tongue. “I” can taste a grape and I know a computer can't do that. Again this is obviously true but, and this is vitally important, when I see “you” taste a grape “I” don't experience that taste only you do !

I can look at you and “infer” the taste you are experiencing but how do I actually know it? - Only by observation. I can only figure out if you like the taste of that grape by hearing you say

“mmm” or make a smiling face and reach for another one but that's all I can do - and a computer monitoring your output after tasting that grape has all that same data that I have - and can process that data faster and more accurately than I can by far.

To a computer, observing you taste a grape or hearing a song you like is the same thing – its just an event that causes you to react and that reaction creates data. Interpreting that reaction that you “output” is also the exact same data that “I” use to decide if you liked that grape (or song).

The computer will never be human and actually “taste” a grape but it doesn't have to - it can make the exact same conclusions as you would because it relies on the “exactly” same data!

Since the observing computer can gather far more data than you can it could even be argued that the computer will be far more exact in its judgments. This is the opening of a door for the computer to make many more far ranging judgments in all related categories.

Since that same computer has a stored and immediately accessible database of all medical and nutritional knowledge it could easily decide if that grape is good for you or compute the optimal number of grapes you should eat. It knows if you recently expended a lot of energy and are, at this moment, deficit in certain vitamins that a grape contains.

You can see where this goes and at light speed. Let's go a little further down this rabbit hole.

We now agree that the computer can decide as well as we humans can if you like the grape you just ate – and it also knows everything in its database which includes all the science on the various chemicals in grapes.

Now, let's get that computer to augment the next grape you eat with some chemical changes and then gather all the data generated by your reactions after eating that grape. After doing that experiment a million times using feedback loops it will be an easy computation to decide which grape tasted best to you. Grape sellers will rejoice because they can now design “new” grapes that are more desired than the “old” grapes.

We could call this just a market survey on various grape varieties but it's also a way to program you to like certain grapes. Just change the taste to whatever it is that your responses indicate that you really liked and you will now buy these new grapes. You are now being manipulated by adjusting that “input” to program the desired “output”.

This concept applies generally as well. With the ability to alter the “input” you can create the “output” desired and you can perfect this procedure with millions of simulation runs. Every parameter of that grape can be tested and tracked across all of humanity very quickly. The manipulation possible here is extensive.

A very beneficial use of this newly created ability of the computer would be to enhance the search for cures for cancer by assembling the entire database of inputs to humans and isolating which of those humans get cancer - the GPUs can then look to find patterns in the inputs. This of course, is exactly what current cancer researchers are doing but the computer can do it faster and better isolate interesting patterns from among the trillions of possible combinations. Humans doing this research, are

trying to get lucky in finding those combinations - the computer is just able to crunch far more complete and far more complex data. Since cancer is not yet cured it's obvious that nobody has yet been able to see those hidden patterns. In this category this manipulation of variables by the computer could be seen as a great option for humans.

In the music world this manipulation is already taking place - it's hidden in "market surveys". Major music companies already know which music sells to the masses but as data collection becomes more complete and can be linked to each individual person, they will soon know exactly what your personal tastes are and you may be able to order a song you like that was personally prepared for you - and maybe also some grapes that taste great.

Of course "grapes" as presented here are rather silly examples but the ability to create a desired "output" by adjusting the "input" is a slippery slope. Google adjusts "input" to you so the output is that you buy that light from Home Depot - CNN and Fox both adjust what you see to try and influence your political opinions. Every advertiser is trying to do the same - the "input" is their ads and the "output" is hopefully that you buy their product.

As data collection continues these "input" adjustments will become better and better. Ads are typically aimed at larger groups to try and stimulate sales but are increasingly able to target individuals. Ever told a friend via email that you were thinking of taking a trip to the Caribbean - and then magically you have pop-ups from travel agencies ?

All the search engines currently tailor your search results to you "personally" based on your personal data they have been collecting . This collected data is then sold to anyone that wants to sell you something or maybe adjust your thinking. We do have some privacy laws but does anyone believe they are honored? I think it is obvious that they are not.

Cameras everywhere, phones in every pocket, satellites, smart meters, Siri, facial recognition etc. - you can't move very far in any direction without creating collectible data. We are each massive troves of data literally down to our individual DNA all of which is already collected or soon to be collected.

As we mentioned before there is no reason to spend the massive money on data centers other than to collect and store this data. Being able to monitor all human "outputs" together with the ability to amend all the "inputs" is a dystopian amount of power albeit with some very beneficial uses as discussed .

Along the way we have noted some seminal events that inform us as to the potential of AI. The biggest single event was the processing abilities of digital computers empowering the exploding speed of computation but also the additional abilities created by GPUs. Without this new category of chips current AI technology would not be possible. The massive data we are collecting could not be accessed quickly enough to be useful - nor could complex patterns be recognized which lies at the heart of intelligence.

The training data for AI is essential everything any human has ever done or denoted including all books and all events that were ever witnessed and any ideas that were ever communicated. Current AI models like ChatGPT still make some errors due to the fact that up until now they have only been

trained with data from previous human actions. Humans make errors and by relying on this imperfect information errors are assured. The new tech now being employed is categorically different. When ChatGPT responds to you your reaction itself becomes a data point. The massive collection of your responses is the new tool that can reshape the taste of that grape by adjusting the inputs to achieve the desired output. This feedback loop is the “self-correction” mechanism that was missing until lately. The collection of this data is ramping up now and the GPUs necessary to process it are being assembled as is the electrical energy sources necessary to power it.

Also remember the AlphaGo/Master progression - AlphaGo used pre-training from past human actions - Master entirely trained itself without knowing any human history. We don't yet have a handle on where this can go. This is at the heart of the real AI unknowns at this moment in time.

Upward and Wider

Many of the ideas about AI capabilities contemplated herein rely on massive data collection which is not yet fully implemented but, that collection is increasing quickly as is the ability to process and organize this data in usable form. AI in general will be a huge change in setting the direction of humanity - far larger than the effect of the internet which will in the future be seen as just the first showcase of digital technology.

During the entire history of life on this planet “natural evolution” was the method that kept pace, adapting to the constantly changing environment. This was true long before humans were here and actually continues to this day and it obviously worked very well since we are here in great numbers. Currently, with the overwhelming amount of humans here and the rapidly advancing affect we are having on the biosphere the evolutionary pace that has always worked is now far too slow. If we are to continue existing something else will have to take over - and that something will likely be some version of AI.

We will have to use our thinking minds to develop strategies and new ideas to accommodate as yet unseen events

The developing AI technology is a very interesting topic to contemplate and if considered on its own, it is easy to consider this as just a field of inquiry and discuss it as a discrete topic. However, the very nature of intelligence is not a separate field. It is the guiding force of humans on Earth and cannot be considered as separate from our history or our future.

We have determined so far that both huge data collection abilities and the compute power to manipulate that data are necessary assets to have to enter the AI arena.

With this in mind, I think it's imperative to discuss “who” has access to collecting this data and “who” has the compute power since this will shape the goals of the technology.

Whose Show is it ?

At this point in the US, the only groups that are able to collect this data are the Big Tech

companies. They have developed the methods and storage technology necessary and they have funded this procedure by monetizing the use of this data. Apple, Google, Microsoft, Amazon, Facebook, Meta, Open AI — it's very obvious that all of them can and do collect every piece of data they can access since this is the pre-training data necessary for advanced AI. It's no surprise then that all the current AI players are either these companies themselves or are directly funded and partially owned by them.

ChatGPT is OpenAI - Amazon, Bedrock, and Nova are partially owned and funded by Microsoft - Gemini is Google DeepSeek - Meta is Facebook - Grok is Elon Musk, who was also involved in OpenAI. These are the only real players capable of collecting the necessary data for AI models. They all necessarily also have business and monetary connections to the hardware suppliers like NVIDIA, AMD, and Intel for the actual chips, and Google, Amazon, HP, and Microsoft for peripheral services like cloud services. It's all an interconnected network and you ain't in it !

You can invest in it via the stock market, but you are not driving the boat in any meaningful way and most of the important technologies are industry secrets and protected by all means available.

Of course, these tech companies cannot operate without Government rules and laws allowing them to thrive. So it's inevitable that Big Tech will seek to influence those laws. Big money always speaks in those arenas, but so does access to “big data”.

This “big data” can shape public opinion and, via the methods we have been discussing, can also influence the direction of those opinions. Any politician that needs to be re-elected wants this manipulation service – and that can only be implemented by the tech companies.

A “Duopoly” of power is thus formed by Big Tech and Government, both needing each other for their continued existence. Tech companies need favorable Government laws and regulations to ensure their monopolies on data collection, and the political class needs the big data to craft re-election strategies to keep them in power. It's an incredibly resilient structure that is fully in place right now.

Look at the money flowing to the political class via both donations to campaigns and also likely in under-the-table transactions that are more difficult to spot. Conversely, there are NO effective laws or regulations on the entire field of AI that substantially slows their progress. All obstacles to full speed ahead have been removed.

This ability of AI to change and guide public perception is very evident in the discussions, or lack of same, about how to “power” the data centers. A number of years ago there was extensive public discussion about the safety of nuclear power plants both here in the US and Europe. The worry over safety basically stopped all nuclear power plants from being built and in Europe caused the shutdown of most operating plants. This was a “public” discussion that changed the course of power generation strategies for many years. The important point to take notice of here is that these discussions were public and that they were a decisive component in the policies moving forward.

Today, many of the AI data centers are currently recommissioning closed nuclear plants or building new ones to supply their huge power needs. The old nuke plants that were shut down years ago in Three Mile Island Pennsylvania, Palisades Michigan, Duane Arnold Iowa, and Susquehanna Pennsylvania are being recommissioned and are being financed by Microsoft, Amazon, and Google.

The Nuclear Regulatory Commission has changed the multi-year process for new plant building licenses down to 18 months and has at least 10 plants planned for opening before 2030.

The point here is not whether we should be doing this or whether it is safe, but the obvious fact that there has been exactly NO public debate - and virtually none of this has been reported or questioned by any major news agencies. These recommissioned plants are already approved and en-route to opening.

I think it's clear that the shaping of public awareness is involved in the total lack of public discussion on this matter, demonstrating that the thought-shaping technology is operating efficiently. Both the Tech companies and the Government have vested interests in these matters and have been almost totally effective in subverting public discourse.

Once these public control procedures can be seen operating in venues as large and important as these, we must assume that they will become ubiquitous, and as their data and compute assets increase, will become formative elements in guiding our opinions and by extension, our future.

Duopoly

I think this “Duopoly” of control is going to be the most important construct for the near future, and that future is starting right now. Do we want this union of Big Tech and Government to have sole control of the ability to shape human opinions into the directions that will enhance their own power? That does not sound like a good idea.

Totalitarian governments don't separate Big Tech and Government since tech is conjoined with, allowed and guided by Government. They can operate quickly and efficiently in this arena without public debate.

I think it's clear that our Duopoly of Government and Big Tech is actually operating in this same manner, although embedded in a different governmental structure and employing different procedures to attain their desired goals.

Is our voting procedure functioning as it should? It's all wrapped up in computer procedures to effectively count the vast numbers of individual votes. The only people who have the technical abilities to oversee this procedure are the very same people who stand to benefit from the results. The laws governing all these procedures are also made by the same group who benefit from the results - a design that almost always assures corruption.

As AI permeates life in general and continuously gains power, speed, and the demonstrated abilities to shape human awareness, we had better decide who controls this technology.

I have no idea how to do this. Whoever controls this technology will have unprecedented abilities in the human arena, abilities never before available and at a scale never before seen in all of human history. If we do find a way to create some oversight in this arena, who is going to oversee them? And who will oversee the overseers?

AI is coming for sure. Some will try to stop its development, but as long as anyone on earth is working on it, it would be a foolhardy move to stop our own progress. Given the control abilities we have already seen and understanding some of the actual technology that will be used to implement these advancements, it's clear that it is coming fast and will reshape future mankind one way or the other.

Vexing problems for sure, and a great and existential dilemma to be left to the next generation to solve.

Some Computer History

The history of computers shows that they have been moving up the chain of complexity in the tasks assigned to them. First came the calculator, then word processors with advancing abilities like cut and paste, then the internet, robotics, display technology, and then penetration into the arts like music, which we have detailed a bit here. They have moved up the chain of complex tasks as the computer tech has advanced in speed, miniaturization, and cost, and with AI has reached the “pinnacle” of complexity, trying to do everything a human can do.

This could be considered the modern Turing test on steroids and once again will call to mind the Van Gogh effect. The computer doesn't need to be a perfect replica of a human, only close enough so we can't tell the difference.

So far in this historical timeline, it is obvious that in each category the computer vastly outperforms any human doing the same tasks faster and more accurately. Suno can create music that is at least difficult to tell if it was written or played by a human. Photos and videos are now rampant that are on par with this and improving daily. Bots are omnipresent everywhere and posting content that is difficult, if not impossible, to identify.

These observations are the visible and audible evidence that can provide a tool to measure how close AI is doing in its current abilities. I would say at least pretty amazing already.

Meet Andy

Since we are trying to get a handle on where this can go, I would like to suggest a mental thought exercise that may make our guesses better. Meet “Andy the Android”. He looks like a real human. He can talk, listen, smile, and walk exactly like you can. Of course, today's Andy is not quite to this level, but that is only a matter of time.

Previously, we have wondered how a computer could possibly taste a grape - and probably balked at the idea that your laptop could accomplish that task. I think it's our error that we think of it this way.

Andy is a computer, isn't he - just housed in a computer case shaped like a human. It's much easier to imagine Andy eating that grape and, as we have previously discussed, displaying all the output indicating that he likes the grape. If Andy really looks indistinguishable from a real human, you have

absolutely no way to tell that he (it) isn't tasting that grape.

Let's see if we can trip up Andy by experiment.

We look up studies done on grape appreciation and find that 67% of people like grapes. Now, we line up 100 “Andys” and feed them all grapes and watch the reactions. If Andy is programmed to display the “I like grapes” output, we should see all 100 Andys do that, proving that Andy is likely a programmed android.

However, the Android computer already knows that data you looked up and simply adjusts its output so only 67% produce the “I like grapes” output. Andy's memory banks have way more information than you could ever have and can adjust its output instantly. In fact, the computer can go much further by finding out other patterns in chemistry, digestion, and energy output that mammals have that seem to precede hunger and that scientifically predict the 67% measure that you looked up.

The computer doesn't know if it is housed in a laptop case or Andy's body. Both are just operating according to an algorithm or a program responding to various “sensor” inputs. So - in addition to an advanced computer, Andy needs some “sensor” inputs relating to our human senses of seeing, hearing, touch, and smell.

Seeing is easy. Video technology can already read more visual data than our eyes and with vastly better detail. Just because the light spectrum our eyes process is between 380 (violet) and 750 (red) nanometers, there is no reason Andy's eyes need to have these same limits. Night vision goggles are made to see the longer wavelengths below red and Andy's visual sensors could be seeing this easily with current technology. This is sort of the reverse Van Gogh effect, where Andy's visual sensors read incoming visual data far better than our human eyes.

Hearing is also easy. We have already discussed digital encoding of real-world analog sound. Current technology is already capable of easily fooling your ears, both in the production of sound and also in analyzing audio inputs, again far superior to our old-fashioned ears.

Touch and smell are still a little behind the curve here, but still making fast progress. Smell should be fairly easy since it's really just chemistry and should be more easily modeled, especially since Andy's nose is connected with every chemical sensor ever invented.

Touch is a tougher challenge since we have entire bodies as input sensors and because it is often highly correlated to sight. Data gloves, VR (virtual reality), and both industrial and sex robots are all pushing development in this arena with lots of products already commercially available.

Andy doesn't need to “experience” any of this input. He just needs very accurate sensors responding to the real world and sending that data to his computer so he can recognize what “output” to produce.

Remember the Van Gogh example - once the dots reproducing the copy of the real painting are small enough that we can't see them, the copy functions in all ways equally to the real painting. Same with Andy - once his “outputs” are so good that we can no longer tell he is an android, to us he is functionally a human. Andy now passes the Van Gogh test!

Andy can't “feel” being hungry, but he can see it as a pattern. He can see all humans eat, how much they eat, and when they eat. By monitoring speech, he will know that the first meal of the day is called “breakfast” with bacon and eggs more often for breakfast than dinner, all easily noted as recurrent patterns. Your cat gets hungry too, and the computer can see some parallel patterns there and with speech monitoring begin to understand that all mammals eat.

In its pre-trained memory banks, the computer has already digested all encyclopedias and medical literature on mammals and hunger and cats. This data bank is huge, but with incredible computer speed and GPUs locating recurrent patterns, the computer can easily cause Andy to mimic the output that we humans do that, for example, indicate we are “hungry” or even enjoyed that grape....

The computer knows all of history and adds all currently generated observations to build its database. With all of this information, it can extrapolate patterns and generalize further to understand even how hunter-gatherer societies planned their days around obtaining food and how and why the first agrarian communities formed.

You can see how easily patterns can get linked just by watching humans get hungry and identifying the actions those patterns cause. Andy can't experience hunger, but he can link those easily visible patterns and perfectly mimic them and if he can do that, how can we tell he is an android and not a human?

A Science Diversion

In order to enhance our thinking about how far Andy can go, I would like to introduce my old friend Albert Einstein.

At some point in his life he was watching something fall and wondering if there was anything more to learn about that event. Isaac Newton already had a concise understanding of gravity and had the math to prove it. His laws of universal gravitation were fully accepted by the entire physics community and had been since 1687 and those laws were tested and proven over and over using all the methods available at the time. They predicted how apples fall and at what speed, and how the planets should move around the sun. A couple of small anomalies to these laws were also known by 1905 when Einstein proposed the Special Theory of Relativity and a final General version in 1915.

It was known, for example, that heavy objects fall at the same speed as light objects - Newton's laws could not explain that. Astronomers at this time also noticed that Mercury's elliptical orbit was slowly changing over time when Newton's laws predicted there should be a stable orbit. The change was very, very small and only observable over hundreds of years but still unexplained.

Einstein proposed a “thought experiment”. Imagine a box closed on all sides with you inside that box. Let's make it look like an elevator with you standing inside, stopped at the tenth floor of a building. You decide to do some measurements while in the box - so you weigh yourself and denote the number. You also drop an apple and measure its speed as it falls.

Now imagine that same elevator box being in outer space far away from any planet. Attached

to the bottom of that elevator is a rocket motor with an adjustable thrust motor. You turn up that thrust knob until the acceleration matches the speed that you measured when the apple fell back on Earth in that elevator stopped on the tenth floor.

Now the question - can you do any physics experiments in your outer space elevator box that would yield a different result than when you did the same exact experiment standing still on the earthbound 10th floor box?

Drop the apple in your outer space box and it goes to the floor at exactly the same speed as in the earthbound box. Of course, if you were standing outside that box in space, you would see the apple standing still and the elevator box accelerating upwards - but from inside the box, it would look and feel identical to the Earth Elevator experiment. As long as that rocket maintained the acceleration you adjusted, your body would also be affected the same way as the apple. You could weigh yourself, pinned to the bottom of that box in space, and you would measure the exact same weight as you did on Earth.

In fact, there is no experiment you can do inside either of those two boxes to indicate which one you are in !

Let's reverse this test by stopping the rocket motor. Now you and the box are floating in space. Drop the apple, and it stays by your hand - you and the apple are floating weightlessly. What would you have to do in the earthbound elevator to mimic that weightlessness? You would have to cut the cable and let that elevator free fall. You and the apple are now weightless, at least until you hit the ground floor.

From these tests we have determined that “acceleration” provided by the rocket motor is indistinguishable from the “gravity” you experience here on Earth. Einstein called this the “Principle of Equivalence”, and this, together with the discovery that the speed of light is a constant (along with some imagination and some math skill) yielded the idea proposed in the General Theory of Relativity that space and time are not fixed or separate quantities but always linked in “spacetime”

Gravity is not really a force - mass creates a warp in space-time and we experience that warp as gravity. By using Einstein's new theory, the precession of Mercury's orbit was explained and the observation that all objects fall at the same speed was now completely understood.

It's now obvious that when you are in the free-floating box in space with the rocket motor turned off, any 2 objects you let go of will float, unmoving regardless of their weight and if you fire up that rocket, both of those objects will get to the floor at the same time. Easy to see and understand now - that all objects fall at the same speed regardless of their weight. Einstein's theory also predicted that time will pass at a slower rate the faster you are moving!

Newton's laws were judged to be perfectly correct for 220 years, partly because they are a very accurate representation of what we humans can observe. All human experiences were logged here on Earth on a planet with a steady mass and moving smoothly through space.

You can't see any time dilation until you are traveling at a good portion of the speed of light (186,000 miles per second). Newton couldn't go this fast. We now have satellites that go thousands

of miles per hour and do experience time dilation. Even though that speed is a very small portion of the speed of light, it must be accounted for to keep our GPS units working correctly. The clocks in the satellites must be perfectly in sync with the GPS clocks in our cars.

The clocks in the satellites run slower than in our cars because they are moving faster than we are, and the GPS clocks in our cars move slower because we are in a higher gravitational field. For our GPS units to remain accurate, this must be continuously adjusted and was perfectly predicted by Einstein.

Linking acceleration to gravity and linking space and time seems counterintuitive, but doing so properly explains why clocks move slower the faster they travel and that space seems to compress along the direction of travel - and even why gravity affects time.

Newton assumed space and time were separate and continuous - he assumed space is unchanging throughout the universe, and that time passes at the same inevitable rate everywhere.

These seemingly common sense ideas were completely wrong.

Your twin sister gets on a speedy rocket ship and flies away, returning 10 years later. She steps out of the ship on her return and she is only six years older. She experienced six years of time while you experienced 10 years. Einstein explained this both conceptually and mathematically.

Amazing, strange and “seemingly impossible” sounding events are suddenly not only possible but inevitable. With this new framework of thought, soon astronomers could see that the entire universe was getting larger and actually measure the rate of expansion. By extrapolating Einstein's equations, it was found that the planets, stars, and interstellar dust were only about 5% of the mass necessary to account for this expansion rate. Dark matter and dark energy account for the rest.

It seems crazy to think that everything we can see in the entire universe accounts for only 5% of what apparently exists. Dark matter at 25% and dark energy at 70% account for the rest. Not only can't we see either, we currently have no idea what they could possibly be made of.

By adding the famous $E=Mc^2$ formula, we can now see that mass (weight), and energy are essentially two forms of the same thing. The early universe was likely all energy, some of which converted to mass as it expanded and cooled, another equivalence that has cosmic consequences - and by extending these equations even further, we find black holes.

So, how does all this relate to our discussion of the future of AI?

In one day and with one set of new equations in 1915, Einstein changed and upended 220 years of what we had considered “unquestionable” truth.

Things as basic as “time” and “space” were no longer fixed commodities but now dynamical systems, and they spawned entire new arenas of thought that were unimaginable before this time.

Our “intuition” that time just passes steadily at a fixed rate and our idea that it will be the same on a speeding rocket was wrong. Likewise, we see space as a fixed commodity - we measure the

distance to a star to be 10 light years, but when your twin travels there in a speedy rocket, she measures the distance as only 6 light years.

It seems totally impossible, but “both of you are correct”. Once again, our “intuition” is wrong.

We build and learn to trust our intuition here on Earth, and in that frame of reference it works quite well - but when we expand our reference frame to the size of the universe, we must learn to abandon those intuitions or they will mislead us.

The same occurs when we “miniaturize” our reference frame to the size of atoms - the quantum world is even more non-intuitive. At that size, the world is not just like a smaller version of the world we see and live in. It is actually an entirely different place with different laws and rules. Quantum indeterminacy, wave-particle duality, Heisenberg uncertainty principle, Hilbert space - a whole new lingo is necessary to explain this arena.

This is the world in which your computer lives and operates.

An electron is a quantum particle, which is not really a particle but a wave, at least until you look at it - then it's a particle. Where is that particle? Its location is explained by Schrödinger's wave equations, which cannot specify where it is, only percentage chances on where it might be. That quantum particle lives in a “superposition” of all places it might be.

This is not like our daily world where we know for sure where that baseball is. It has a definite position and momentum, but it turns out our intuition is wrong again.

The location of that baseball is, in reality, governed by the locations of the trillions of atoms making up that baseball, and each of these atoms is only locatable statistically. Because that baseball is “large”, it seems to us to have an actual location, but at the atomic level there is no specific location. That baseball is actually a fuzzy sphere moving through the air. The fuzz is made up of all those quantum particles moving around statistically in their superposition locations. What we see as a baseball is actually only an approximation of all those trillions of quantum particles moving around. The fuzz is way too small to see, but it is there. In the same way, your hand and your baseball glove is fuzzy too.

This bizarre excursion into Einstein's relativity and quantum mechanics is pertinent to our thinking about AI. In each field, our notions of reality are stretched in unimaginable directions. We should keep this in mind when guessing at where AI may go and how it may advance.

If Einstein could show that seemingly unrelated phenomena like acceleration is somehow equivalent to gravity, is it somehow that much different to equate Andy's ability to perfectly mimic human responses with being human? If we can't tell the difference, how do we know?

Einstein's revelations about space and time and the entirety of quantum mechanics all existed when cavemen populated the earth - we are only seeing a closer look as our understanding advances. This may be a better perspective on AI as well. We may just be beginning to get a clearer look at what reality actually is.

Andy again

In Andy's case, he can easily see the same world we see, his database has all of recorded history and his sensors take in current events as a data stream. With just this, as we have just seen he can easily adjust his “output” to perfectly mirror human actions and reactions.

We humans “feel” gravity and think we understand “space”, but as we have seen our built-up intuitions are actually wrong. If Andy's database contains all Einstein's equations and all the quantum mechanics specifications, he will not have to fight any “intuitions” - he will just see all of it as a coherent whole - with the erroneous conclusions that humans have made based on a severely limited perspective and actually limited by our intuitions.

This should radically change our judgment of Andy's range of outputs.

With the computer speeds and parallel processing already existing, his possible outputs are actually unpredictable. Remember the AlphaGo computer that beat the existing human champ, and then the Master that easily beat AlphaGo. Master was totally self-taught and was making moves that no human had ever thought of before and was unrecognizable even to the programmers.

This was 10 years ago and of course Andy already has this in his database, and his computational apparatus has expanded exponentially as well. We don't really have a clue as to what Andy can do now, and even less of a clue on tomorrow's Andy !

Acting like a human is easy for Andy using only a small subset of the data he has available. We need to be thinking in grander fashion about where he can go tomorrow.

Thinking of the computer as “Andy the Android” has helped us to expand our thoughts, but of course this is very limiting as well. We humans have our databanks filled with our experiences like tasting that grape. Andy doesn't have that particular experience, but instead has unimaginable amounts of data surrounding “eating” and “grapes” and “chemistry” as well as every observed action that humans exhibit when eating said grape.

His data goes all the way down to the quantum level and all the way up to Einstein's cosmology. That Andy will likely come up with some AlphaGo/Master moves that no human has ever done - easily. It's inevitable.

I'm one of those humans, so I should be aware enough to realize that Andy's outputs probably can't be predicted. We can throttle down his outputs by applying programming that tries to limit his actions to our specified parameters, but we may be unaware of where those limits are and how we can describe and program them. There are a lot of people working on this “tuning” problem, but it's still uncertain how and if we can operate in this arena. Good luck to us!

Just remember that poor guy in outer space, in the box with a rocket attached. There was absolutely no way he could know if he was in that box or sitting on the tenth floor of a building here on Earth. Not only could he not tell, but there is no experiment he could do that would tell him.

We should remember this when we prognosticate on where AI can go. We have no tests, and even if we did, they might only be valid in our little box.

Another Thought Experiment Please

Einstein's equations show that the passage of time slows as your velocity increases, but this slowing is not proportional. At half the speed of light, your time clocks are only about 15% slower, and at 75% of the speed of light they are still only 51% slower. His equations also show that the faster you go, the energy required to increase your speed grows exponentially, so ever reaching those speeds may be unlikely.

Our Milky Way is about 150,000 light years across and even at half the speed of light it would still take 75,000 years to cross. The nearest galaxy to ours is Andromeda and it is about 2 million light years away. If you do the math, we can guess that no “human” will ever be able to do those trips.

If we are ever to really begin to explore the universe, we will have to send Andy to do it. We are already doing this a bit with our little gizmos we send to Mars - they are actually baby computers. Soon we will be sending Andy. He doesn't age or need to eat, and as long as we feed him power he can travel these distances.

So - our thought experiment is this. Imagine a UFO lands in Washington. Out steps something. We know that no other human life form could possibly have gotten here due to the limits of physics so we know it is some form of computer - some other version of Andy.

Now let's look at “us humans” through the eyes of that visiting computer.

It would see a head with some strange meaty substance inside that seems to guide our displayed intelligence. We would look incredibly foreign to that alien computer, and they would have no idea “how” our computational skills are created.

Seems to me that describes” exactly” what we see looking at our advanced Andy - a device that produces intelligent acts where we have no idea how those acts are created.

We are in a box accelerating through space with no way to tell anything more.

This is a good description of where we are now, gazing at our AI attempts. The range of possibilities is far beyond our ability to guess from inside our little box.

This scenario is categorically different from anything in human history.

The wheel and the printing press, and even the computer, certainly expanded human horizons, but they were all created by us, and their effects were observable and predictable. The wheel went where we pointed it, and the printing press only printed what we told it to - Andy, however, can learn and program himself. His eyes are not limited to only the visible spectrum and his ears can register frequencies far beyond what we can hear. His internal database contains everything mankind has learned throughout its entire history and is updated moment to moment. This data is never forgotten

and is available in its entirety nearly instantly. Pattern recognition far beyond what a human could accomplish is spread over this vast database and is unhindered by our human limitations and will certainly yield some unexpected results. That Andy certainly is an interesting character.

Acceleration is equivalent to gravity. Electricity and magnetism are two incarnations of the same thing. Photons have no mass, and yet they exist and are the carriers of electromagnetism. The universe may have been born at a moment where its entirety was comprised of a mathematical point of infinite energy.

For us, it is a big stretch to try and really understand any of this. For Andy, it's not intimidating at all - it's just a vast playfield of data that may contain patterns as yet undiscovered.

Let's be sure to remember Andy is still actually a computer. We built him in a human form to look like us and be able to gather data just like we do, in order to more perfectly mimic what we do - but Andy is still just a collection of sensors sending data to his processor. It could just as well be your Laptop connected to these sensors, but it's easier for us to visualize and interact with if it's shaped like a human.

In order to be AI, this laptop needs to be connected to all this data we have mentioned in order to do the amazing things we expect Andy to do. Let's examine this a bit closer.

Back Into the Mud

When you Google on your phone to find out some facts, you are swiftly shown that data. It was picked up from whatever source was accessible, likely through the internet or in some cases retrieved from some stored databank. Every part of this process requires a lot of infrastructure and compute power to accomplish and was very expensive and labor-intensive to create.

In simple terms, the Big Tech companies spent the money and provided the expertise to do this in return for the huge money they could make by selling the data that they could collect. That data is valuable as a product to sell but even more so when used as a means of control. We have touched a bit on this previously, but it bears repeating that the more data they can collect, the more money and power they accumulate. In the AI game, this data is the pre-training for Andy.

They also need big computer power to manage and process this data, so it's no wonder that all these Big Tech groups are also either computer companies or intimately connected to them. With the "huge" leap ahead that GPU chips have provided in AI, it's also apparent that these groups need "huge" electrical power and access to water for cooling in order to operate these mammoth planned data centers.

In order to obtain these necessities, pairing with Government is essential. Power and water at this scale require public infrastructure spending that is the province of Government. The money generated by the AI companies is like lubrication to the political system and is dispersed in huge amounts to those in Government responsible for the laws and statutes that regulate these categories. The Tech companies are big businesses and need to see a return on investment for those donations and receive it with the political favors necessary for them to flourish and maintain their monopoly.

Everyone in the political class also needs to be re-elected regularly. Money certainly helps here, but even more valuable is the data that is being collected continuously. I have touched on these facts previously, but I bring it up again because I think this is currently the biggest dilemma that needs to be addressed. This is the elephant that is in our room right now.

You cannot let any single private, uninspected entity have total control over what humanity is allowed to see. It will not work well. The depth of control in all the ways we have tried to explain herein is growing by the moment. Data centers are popping up unquestioned everywhere, illustrating the necessary intimate connection with the Government is functioning unimpeded. Since I started this mini-book, Oracle has released plans for its Solstice project running 100,000 NVIDIA GPU chips.

The obvious control of media, internet, and education is advancing from an already intrusive amount to another level entirely. It's already very difficult to find trustworthy sources of information, and with the AI skills already implemented, it will become continuously harder.

We can still use a bit of intuition to guide our truth-seeking, but I think it's very clear from some of the examples we have studied that intuition is often wrong and also easily misled. Once our training data is corrupted, the conclusions formed must also be suspect.

I really have no idea how to control or even regulate this runaway train - AI is coming whether we want it or not. It can certainly be presented as a force for good, and it's likely some great advances will happen, but the problem of allowing control to be in private hands will remain.

To date, there has been no effective public debate on any aspects of AI—not the energy use, the intrusive data collection, or the obvious connection between Government and Big Tech.

Events yet to unfold will certainly affect the trajectory and speed of AI rollout. Discovery of new GPU chips that use much less electricity will certainly leave a lot of investors with obsolete chips, but others will step in to pick up the pieces.

There will, of course, be big scrambling in finance and Government because big money is involved. This will certainly affect the investors and Government players, but AI progress will continue due to this big money generation and because of the social control that comes along with it. Politically and monetarily, some may win and some may lose, but nobody will sacrifice the social control.

Some companies will waste too much money chasing down a dead-end idea and go bankrupt but the winners will always be favored by Government for the obvious benefit of both parties. As we speak, data centers are being built, new plants are coming online, and massive money is being spent with basically no oversight.

The undiscussed risk is that data centers need the same resources as humans, “power and water”, and the data centers all have long-term contracts for those necessities.

It is often described and marketed as a “public-private” partnership, which sounds good, but the “public” is the Government and the “private” is about six or seven people who have become the richest

people in the history of the world. Both parties have total control of a structure that they both profit from and are keeping it that way.

Since 2000, computers are everywhere. Emails, texts, phones, vacuum cleaners, cameras, satellites, cars, doorbells, on your wrist and in your body - this is part of the network empowering the massive data collection infrastructure and is purposely invisible to us behind an opaque screen.

If you look at history, many of our advanced electronics and computer upgrades come as spin-offs from military use. The Internet is a great example of this, as was Google, both of which received initial funding from the Government. This created the Government / Big Tech connection right from the beginning, and both benefit far more if it is not watched.

I have previously called this Government–Big Tech union a “Duopoly”, which is actually the union of two monopolies.

It's time to recall what “Techno-Authoritarianism” is - a society where the Government and mega corporations merge into a single directed Duopoly that has constant real-time access to all data for everyone under their control. With the speedy rollout of data centers and massive compute power, this gives them the power to monitor and control every transaction by manipulating access to information.

In this document I have often referred to this collected information as “input”. By controlling all “input” to people and monitoring all the “output” from those people in real time, the control mechanism is easy to see, especially when you realize that the monitoring can be done nearly instantly and humanity-wide.

Societal tuning has already progressed far beyond CNN vs. Fox and is often quite easy to see. There are bots targeting groups and individuals all across the internet, and most media is already aligned with visible Government influences, as are education programs and design. Of course, adjusting education can be seen as “tuning the future”. The timeline for societal control can easily span generations.

AI is making all this manipulation much more effective by monitoring “outputs” and then simply adjusting “inputs”, and can be done in real time. One question that should be asked is, “how valid is the concept of “voting” when the public mind is so easily tuned by the Government / Tech Duopoly”?

Our little individual voices can still be heard, as evidenced by this little mini-book you are reading but the data necessary to make wide-ranging and effective conclusions is increasingly only available from the Duopoly. Any technology which has this power should obviously be watched closely but it can't be - because it is more valuable to each of its stakeholders if it is privately held.

The coming more advanced AI capabilities will intrude in as of yet unnamed ways, literally down to the level of DNA - which is also actively being collected. Cloning for organs, CRISPR adjustments to modify humans' DNA, pacemakers, blood monitoring devices, hearing aids, smartwatches etc, — these are all introductory technologies.

Since AI is trying to create models of how we humans evolve and think, DNA would seem a logical place to look - so let's do that.

DNA and More Complexity

Everything we know of that is alive has DNA. Somehow encoded in that DNA string is the entire building plan for who we are, half from our mothers and half from our fathers, all grown from a single cell.

An adult human has about 37 trillion cells and almost every one of those cells has a copy of that original DNA. That DNA is arranged into a twisted double helix shape that, unwound, would be about 2 meters long. If all those unwound strings in one human body were placed end to end, they would reach from the Sun to Neptune and back four times !

Each DNA string has over 3 billion base pairs along its length that somehow encode all the information necessary to grow and maintain our bodies throughout our entire lifetime.

Current science, believe it or not, has succeeded in sequencing the entire length of the DNA strand by reading the exact order of the joined chemical bases. There are four bases (A, T, G, and C) that combined along that string create our unique biological instruction manual for life.

Since everything we are is somehow encoded here, it is a perfect playing field for AI's advanced pattern recognition skills.

Human researchers are trying to locate these patterns and have had some success, but the complexity is daunting. With 3 billion base pairs and the 4 nucleotides available and slight variations between the 3.8 billion humans on Earth, you can see the dilemma. To add to this complexity, we must also realize that while your DNA is a huge determinant in who you become, there are also the events that happen during your life that affect the outcomes.

These patterns we are trying to locate are not as simple as one base pair means one thing and another base pair means another. There are a few of those simple outcomes, but mostly the interesting combinations involve hundreds if not thousands of interconnected pairs interacting and also not necessarily located next to each other on the 3 billion pair long string. This takes the difficulty of pattern recognition to an entirely new level of complexity and is the reason that current science is struggling with the task.

AI will win this game and here's 5 reasons why.

- 1) Advanced pattern recognition -100,000 GPU chips working with neural networks at blazing speed 24 hours a day
- 2) Access to everyones DNA – not only current humans, but also the accessible historical 117 billion humans who have walked this earth.
- 3) Add in plant and animal DNA for whatever added depth to the patterns this may provide

- 4) Complete history of each person's life experiences, including fully monitored environmental variables, human interactions, foods consumed, societal interactions, and health history through your entire lifetime. This adds to the complexity of the pattern recognition difficulties and is mostly unavailable to current DNA researchers. AI just sees this as data, and it certainly creates a more complex landscape in which to recognize patterns, but it is a necessary component nonetheless. This gives the AI computers access to the absolutely necessary data to link each person's unique health info to their DNA profile.
- 5) Monetary motivation to the AI companies. Easy ones to see would be things like cancer cures, drug compounding. Even the COVID vaccine was the first MRNA vaccine. (RNA is half of the coiled DNA helix).

The only way to successfully navigate this huge data field is by using the computer's ability to speedily handle unbelievably complex variables spread over as of yet unknown ranges.

Human scientists to date are already trying to do this using very diminished data sets. Unfortunately, once again, the only groups having access to this necessarily wide-ranging data are the tech companies, and they are and will continue to monetize access.

Another avenue open only to the tech companies is illustrated by our previous example of the AlphaGo/Master example. If you remember, the Master program was entirely self-taught and handily beat the AlphaGo computer, which had been trained on all previous games. As Demis Hassabis said, "Master made moves no human had ever thought of before". This was over ten years ago and the ability for computers to self-teach has grown exponentially since that time.

In the AlphaGo/Master example, the number of moves all humans had thought of before is huge - but the number of moves no human had thought of before is infinite.

If we are trying to guess what AI can do, we would be wise to remember this example.

Even More Complexity

So now we have, or will soon have, all the data generated by humans containing all the inputs and the related outputs together with the sum total of everything humans have ever learned.

AI will start with this huge data field and try to locate interesting patterns. This is a far greater and more complete data field than any human or even group of humans could possibly peruse. To this will be added the unpredictable skills that can be created by self-training, as in the AlphaGo/Master example.

All patterns we have tried to locate so far through human history are ones that we can recognize. With all the data just mentioned added to the unpredictable results from computer self-training, it's easy to see that the results may be stunningly surprising.

Think of the current search for a cure for cancer. This has already consumed many thousands

of hours of human effort, so far without much success. Whatever patterns coincide to cause cancer have not yet been identified, likely due to their incredible complexity.

It may be the case that the vitamin C in some grapes eaten on a Tuesday, combined with certain pollens inhaled on Wednesday, causes a slight mutation to your digestive lining that gets altered only when you are being jostled about in a moving car on Thursday. This mutation then replicates and causes your pancreatic cancer.

This sounds almost ridiculous, and the linking of those scenarios would likely never be contemplated because of the seeming randomness of them and the daunting complexity of assembling the infinite array of non-intuitive events into the preposterous idea that this particular chain of events can lead to cancer.

Humans would not connect these events, but the computer can and will. This particular chain of events cannot be linked by observing one person or even 100 people. Due to its complexity, the pattern could only emerge by having every data point from every cancer patient over all of recorded history.

A computer trained with all this data would still have a daunting task to find any patterns of interest, even more so when you must consider that the events quoted here will not always lead to cancer. That's why your searchable database has to be so incredibly large.

Current cancer researchers try to collect all pertinent data, but you can easily see that the above scenario would never be constructed simply due to its complexity and because it seems so unlikely. A computer that has access to all that acquired data and then self-trained does not know what is likely or what seems ridiculous. It does not have any preconceived notions. It sees everything as data points.

The computer doesn't say to itself, "everyone eats grapes, that can't be a cause of cancer". The computer just logs the ingestion of the grape and all its complex chemical makeup and adds that event to your personal timeline data, including all your DNA information. The next day when you walk to your car and inhale certain pollens from your neighbor's garden, another huge set of data is logged.

The connection of these events into a provable and testable cause of cancer is so unlikely that no human would spend any time proposing the thesis. The only way any discoverable patterns could be teased out of this array of seemingly random data is with vast arrays of computers working untiringly 24/7, utilizing thousands of GPU chips stacked into neural networks and programmed to consider any interesting patterns no matter where they emerge.

Medical researchers are trying to begin this necessary pattern recognition by listing some foods or environmental situations as "carcinogenic". This is a first baby step to constructing any causal network. Of course eating a food listed as carcinogenic does not mean you will get cancer, but they have apparently connected a lot of data to statistically link its ingestion with increased cancer risks.

When viewed against the incredibly vast array of data points available to the computer network, it's not surprising that we do not currently have a cancer cure. Researchers are certainly using computers more and more in their individual labs, but until all the labs are linked with all the collected data on all the humans and then fed into the huge parallel compute systems, it's unlikely that the

incredibly complex patterns can be identified.

This vast data collection and processing network is not yet fully available, but is speedily being built in bits and pieces with those pieces then arranged into searchable databases consisting of all your personal data scraped from every action you have ever done, all your physical data collected from every test your doctor has ever prescribed, penetrating all the way down to your individual DNA profile.

Somewhere in this vast array of data are patterns that predispose you to getting cancer - or not getting cancer.

The likelihood of any human researcher being able to pick out patterns as complex as we have postulated herein is close to zero. I think this is self-evident since we have no cancer cure, even after thousands of intelligent humans have been looking for any definitive patterns for many, many years.

On the other hand, even with current computer technology, the ability of the vast networks being built to find viable patterns is far greater than zero and over time may approach one.

Some humans get cancer. Somewhere hidden in the vast data created during their lifetimes are some events or chains of events that cause that cancer. No two people have the same chain of events and no two people have the same DNA. Each of us creates a completely unique dataset throughout our lives. That dataset in a searchable database must be compared to every other human's dataset to have any possibility of teasing out any meaningful patterns. The more data collected, the better chances for patterns to emerge.

Some GPU chips will be heating up while trying to locate these hidden patterns.

The Big Jump

The new day that is here is the nexus between exponentially increased data gathering and exponentially increased computer abilities to find increasingly complicated patterns. Currently, the GPU chip is the new technology jump but there will be more as time goes on. So far the GPU chips, which were originally created to handle complicated graphics generation tasks, are being repurposed into the neural networks of today's AI projects. NVIDIA has already invented the H200 chip, which is aimed directly at the AI generation of computers and is even faster than the repurposed GPU chips.

Human experiences and reactions are now logged as data, both historically and in real time. The ability of computers to process this massive data is also growing exponentially, both in the available tech to store and access this data, but also in the massive pattern recognition abilities provided by the revolutionary GPU chips. By combining the computational superiority of computers with the real-life experiences of all of humanity, you get something far superior to either.

People have been using computers for years and yielding many benefits, and of course, humans have been watching other humans forever, also yielding benefits. Now these can be merged in the memory banks of data centers and searched for patterns.

This is the existential jump we are beginning to see.

Where are we now and Where are we going

Ok, let's try to assemble all the elements discussed herein to see where we are now, which will help inform our guesses about where we are going.

I propose the following ten conditions that, taken together, launch an entirely new trajectory for humanity's future. Each of them has been alluded to already but needs to be considered both individually and also cumulatively.

Sometimes human events are additive and sometimes they are multiplicative. A confluence of the following ten points is wildly multiplicative when taken cumulatively. I will list these individually and then discuss them one by one to describe what I believe they bring to the thesis.

- 1) Complete monitoring of events – data collection - “output”
- 2) Ability to create “inputs”
- 3) #1 and #2 together are total control mechanisms
- 4) At this point in time, #1,#2 and #3 are in the hands of Big Tech and Government
- 5) No viable agreed upon mechanism to oversee these functional mechanisms because all inside players benefit from the status quo
- 6) The world is too complicated for any human mind
- 7) Most major guiding events and technologies are already hidden from public view
- 8) GPU’s and Neural Networks
- 9) Data Center compute and storage explosive growth allows #1 and #2
- 10) Computer programming itself (AlphaGo/Master) is unpredictable Wildcard

All of the above are already occurring, but the GPU chips and data center storage growth has amplified the speed and penetration to a completely new level.

Let's cover each of these 10 points individually in greater depth

1- Complete monitoring of events

Can anyone argue that nearly everything we do or say is monitored both legally and secretly? The internet, emails, cameras, vehicles, smartphones, etc. are everywhere. Laws have not kept up with privacy concerns and will always be far behind because the collected data is so valuable both for business and for social control.

The only power sources capable of limiting this data collection are Big Tech and Governments, and both benefit from its collection and implementation. I call this collected data “output” because it is the complete monitoring of your “output” as you live your life.

Not just what you buy from Home Depot, but every word you say, every gesture and facial expression you make, the music you listen to, and how many grapes you eat. This is the collected data

component of AI that is necessary to collect and save. It is expanding and being stored as fast as the data centers can be built.

2- Ability to create “inputs”

This one goes hand in hand with #1. If you can see and hear and monitor all your actions, “outputs”, it's only one more step to feed you “inputs” and then monitor what you do given that input.

This is not new - this is what advertising tries to do - show you an ad that makes you want the product advertised - then try various ads with test groups and monitor the results. This is very slow and very imperfect but still yields measurable results. It is also expensive, and necessarily limited in scope.

AI, however, will be able to monitor millions of ads and see the associated outputs immediately with an essentially infinite number of subjects. This is where the ability to manipulate “inputs” and at the same time monitor “outputs” merges into being social control.

3- #1 and #2 together are total control mechanisms

We have all accepted Home Depot pop-up ads, but this new AI ability is escalating into every niche of human events. Adjust the input to get the desired output and let Big Tech and governments run that show. Neither party wants to let go of that mechanism for obvious reasons.

The collection of data is escalating into every part of life on Earth, all with no meaningful debate on its scope or ultimate use. DNA databases are growing and it's easy to see that the future will likely tend toward more collection of everything.

We have already discussed some of the parameters of DNA research, especially its unbelievable complexity - but as the AI database increases and the raw power of the compute networks also increases, there will likely be some new, unthought-of revelations happening.

I spent a little time in this document discussing Einstein's Relativity and Quantum Mechanics for this very reason.

In both those disciplines, the understanding of this entire world changed, in some cases overnight with a single document or discovery. This is analogous to where we sit with AI.

Data collection (everything you do) paired with enough data storage to have that collected data instantly accessible + enough compute power to process this data and finding patterns far beyond what humanity has been able to discover + distribution networks that allow controlled and tuned access to all this data literally down to each individual human = the things you can see can now be easily controlled, and the things you do are all observed and linked to what you see...

4- #1,#2 and #3 are in the hands of Big Tech and Government

We are currently seeing deep-fakes extend into every field. Music creation, video and audio files, bots on the internet, humanoid robots, AI movie stars, etc. All of these are escalating to where the “real” and the “created” are becoming indistinguishable - we have almost reached the Van Gogh moment.

For all of human history, creativity and intellectual discovery was the exclusive domain of real humans. That era is now over. Not that real humans will not continue to make those discoveries now and then, but we can now no longer tell after the fact if it was actually a human who did it.

The computer does not get tired or need to eat lunch. Its recall of data is larger than all human minds put together, and its access to that data is nearly instantaneous. Its ability to find meaningful patterns also exceeds any group of humans, and its ability to cross-reference any findings is also nearly instantaneous.

It's obvious that whoever controls this set of procedures and the inevitable outputs is well-placed to control humanity,

At this point in time the obvious controllers are the Big Tech and Government Duopoly. They will not let go of this.

5- No viable agreed upon mechanism to oversee these functional mechanisms

As has been mentioned earlier, there is currently no public debate on any possible control or oversight mechanisms. Nuclear plants are recommissioned without debate and in many cases without mention. More and more intrusive data collection continues by monitoring every human action and vast water and electrical resources are harnessed for data centers.

It's easy to see that the Duopoly is already controlling what you can see and read to enough of a degree that none of these large-scale changes are discussed publicly. Once again, the Van Gogh effect is firmly in place. They don't need to control everyone's opinions or minds, only enough so that a public outcry does not occur.

This control was easily visible during the recent COVID event. Anyone disagreeing with the official narrative was silenced by debanking, removals from search engines, and public pressure which was steered by controlling the available data. The “output” of people (their opinions) was adjusted and tuned by the “input” provided to them.

Regardless of your opinion about the COVID event, it is easy to see that the information available to make your personal judgments necessary was heavily controlled both by Big Tech as the implementer and Government as the director and there were certainly hidden agendas driving these events.

Case in point - the requirement that vast amounts of people take the COVID vaccine, regardless of their personal opinions, was incredibly invasive into personal freedoms and body sanctity.

The fact that no meaningful national monitoring of the vaccine's effectiveness or side effects

was done was incredible to me. You have a Government-prescribed shot injected into the largest percentage of people in the history of the Earth with no meaningful follow-up.

You have an incredible chance to test this vaccine with a huge nationwide control group, a first in medical history and you have computers capable of storing all this data, you just need procedures to collect it. Nothing substantial was done.

The VAERS database was optional and difficult to navigate. The continued claims about vaccine safety and effectiveness continually changed over the period, and to date there is still no consensus.

Who made the choice to refrain from meaningful data collection? A lot of money was made by a lot of different entities - did they influence these oversights?

Big Data and Government were responsible for these choices. Big Data controlled the information, and Government made the rules. Not a performance that enhances trust.

6- The world is too complicated for any human mind

A nuclear scientist does not know about virology. A virologist does not understand nuclear reactors. A psychologist likely knows little about either topic, yet we are required to know enough to choose leaders that must understand all these varied topics deeply enough to guide us to making good choices. In turn, they must rely on the current holders of the best knowledge in each field.

Just choosing who those people are is fraught with difficulty as each of them have personal agendas only sometimes aligned with your personal goals. This boils down to “who you trust” and further to how and where did those trustworthy people get their information and is that information correct?

In current times where “inputs” (information) is easily controlled to foster desired actions, it's easy to see that the opinions held by everyone, including experts, can be quite easily manipulated.

We still have all the human manipulation devices such as money, career advancement, envy, power, etc. but increasing amounts of the more subliminal and effective methods are in the hands of the Duopoly and will of course be used for their benefit.

Not only are the individual topics of interest for life on this planet far too complicated for any person to understand deeply enough to make completely trustful conclusions, but at the same time the data can be easily manipulated or withheld.

The stark reality is that all large-scale life-altering topics are far too complicated for us individually to make trustworthy conclusions. We are left to “trust our feelings”, which is about the easiest thing to manipulate and, as experience shows, a very poor guide to the best path.

We do however, regularly need to make judgements – its just a good idea to remember how easily they can be wrong.

7- Most major guiding events and technologies are already hidden from public view

Some things need to be hidden from public view, such as military and weapons research and development, spy technologies, etc. but the current elephant in the room is that data gathering is also mostly secret, both as to how intrusive it is and what it can and will be used for. It's easy to see that this data gathering currently is pervasive in every aspect of life, and the data centers exploding in volume is the endgame of being able to store and use this data.

How can we make valid, trustable conclusions about the Iran War, genocide, nuclear power safety, COVID procedures, DNA manipulation, planetary weather effects, fossil fuels, etc. when much of the pertinent data is either not available or purposely withheld or modified?

Our conclusions can only be as valid as our information. These topics are not only intrinsically very complicated, but when any information is withheld or manipulated, our ability to conclude is further compromised. We are then left to the “trust our feelings” mantra - which is shorthand for “I don't really know”.

Ask almost any protester about the hard facts of their particular interest and it often reduces to their “feelings”. Ask the same of the counter-protester and you get the same answer. Two people vehemently disagreeing about a topic that neither fully understands.

8- GPU's and Neural Networks

We have previously touched on the incredible “parallel” compute capabilities that GPU chips have added to the computer skill set. Neural networks are multipliers of this technology by stacking them into levels, each feeding forward and backward. Data centers are the culmination of this technology by arranging thousands of these networks side by side, all interlinked.

This is the hardware side of AI and is so explosive in its growth that even the insiders and owners don't really have a full handle on its capabilities or even understand its methodology. The AlphaGo/Master self-programming abilities are likely to have expanded beyond our ability to understand and perhaps control.

9- Data Center compute and storage explosive growth allows #1 and #2

#1 was the complete monitoring of events like the “output” and actions of all humans, and #2 was the ability to create and disperse “input” to those various humans.

The data centers are the hardware that is central to these two procedures. They are the storage repositories for the astronomical amount of data being collected and the necessary computers and networks to process this vast database. It is the connections created by monitoring the “outputs” from the various “inputs” that are of interest.

In the past, the inputs to people's lives were simply naturally created by each person's random life experiences. Which book did you read? Where did you travel? Who did you talk to or communicate with? What data did you use to decide to buy that particular lamp from Home Depot?

Increasingly, these “inputs” and “outputs” are being linked by discerning the complex patterns involved in all these choices. The test group is all of humanity, and the vast array of parallel compute-enabled computers are locating these patterns and, in time, likely down to the level of DNA.

This full methodology is not here yet, but is the logical endpoint of all this massive directed compute power left to play with every piece of data generated by humanity.

10- Computer programming itself (AlphaGo/Master) is unpredictable Wildcard

Remember the AlphaGo computer program that beat the current world's champion Go player and the subsequent Master program that beat the AlphaGo computer. “Master” was not programmed by a human. It was given the rules of the Go game and the goal of winning, and then set loose to program itself. According to Demis Hassabis, he had no idea what the process was that created the successful self-programming.

This is a wild card that, by definition is out of our view and apparently out of our control, yet successful at winning the game. Much of the current output of AI is already self-programmed and will likely spread deep and wide across all parts of this technology. Since we don't know how it's doing this self-advancement, we can't really make any valid predictions about where it goes.

Years ago when the pre-trained computer was assigned the task of winning a game of Go and it succeeds, that's an “interesting and surprising” outcome – but when the SELF programmed computer beats the pre-trained computer that's an astonishing outcome.

Today when the task assigned is locating all the hidden patterns governing human actions, the output will also be “astonishing”, especially since this current assigned task has access to nearly unlimited databases and compute power.

Conclusions

Let's see if we can make some conclusions given all this discussion. Referring back to my 10 elements, #1 and #2 are the goal (control) and #8, #9, and #10 are the hardware necessary to accomplish #1 and #2.

In simple terms, once you can monitor all the “output” that a human does and can control all the “inputs” to that human, you have a complete control mechanism. Adjust inputs and monitor outputs completely and at computer speeds, and your control is easy to implement.

This #1 and #2 are first on my list for this reason. You do not have to control every input, just enough to get the desired output. You also do not need to manipulate every specific person, only enough to create the societal movement you want.

This is why and how the “illusion” of freedom can be maintained. There will always be some that are less manipulable or go in different directions. Since you can monitor all the outputs across the entire population, you only need to keep those numbers of dissenters small enough to not redirect the larger group.

Looking at humanity now, we can see the various “methods” used to institute this control. First some effort at re-education - soft power, then some planned societal pressure to get in line and be part of the group, add in some monetary and status incentives and you have a large palette of persuasion techniques. For those that get too far out of line, you have the legal system - just make the dissenting actions illegal and finally, you have removal as a final tool - jail or assassination.

If you are looking closely enough, all of these methods can be seen to have occurred.

Re-education and soft power can and does start with children in schools and continues through life. We can see this control easily by looking at how a society trains their children. Education is “public” and is guided by Governments, which have now merged with Big Tech as the central planners. The control of “input” and the ability to instantly monitor “output” is a foolproof method to guide this plan and all current intellectual disciplines can be impacted.

Psychology is used to design the “inputs”. As psychological knowledge progresses, the input design improves and creates more efficient inputs. This is a feedback loop, increasing effectiveness with each loop.

This, of course, has been done throughout all of human history. In simple terms, this is what “advertising” strives to do: create some input that makes you do the desired output. This can now be implemented society-wide and tested immediately. The best advertising agency is the one that can create the “input” that best causes the desired “output”. Currently, creating the best input and measuring the output is very inefficient and slow. “Input” creation and “output” measurement will soon be immediate, and the entire advertising field may take the same route that recording studios took and become unnecessary.

It's likely that the same thing may happen with the insurance business - insurance relies on monetizing the unknown. You buy health insurance because you have a chance of getting some disease that costs a lot of money to treat. The insurance company calculates the potential cost of your treatment, uses their available data to guess at how likely you are to contract that disease, adds in some profit for themselves, and adjusts your premium accordingly. This data currently makes up the actuarial tables that hopefully capture the percentage risks of various futures.

The entire field of insurance is dependent on these unknowns. As DNA analysis improves, and of course with the omnipresent monitoring of all “inputs” and “outputs”, this field of unknowns diminishes and in some cases is actually controllable. If you are seen visiting the liquor store too often or staying at the bar too late, your insurance premium will go up.

It's already the case that no insurance company will sell you insurance for pre-existing conditions, for obvious reasons. Once everyone is monitored, a lot of the future becomes predictable and the entire insurance business model may disappear.

More Guesses

In this entire document, I have tried to show how seemingly small events or inventions have had huge effects on future life. Often these new advances in knowledge are not even recognized when they occur. Einstein's discovery of relativity and the advancing knowledge of quantum mechanics are major turning points that are illustrated by the vast difference in available technology today versus only 100 years ago.

Electrons are quantum particles and need to be manipulated to create computers and transistors. Your GPS would not work without knowledge of relativity and the distortion of space and time.

It's not necessary to understand Relativity or Quantum Mechanics. You only need to look at today's world and notice the huge changes that these new ideas created. With Relativity, mankind now can look trillions of miles out into space and see back in time to when the entire universe was small and with Quantum Mechanics, we can look down into realms trillions of times smaller than your little finger.

That's an astounding expansion of knowledge within a short time and all this was accomplished with human intelligence alone. To that rapid expansion from just human labor will now be added the explosive computer power being assembled.

More has been discovered in the last 20 years than the previous 80. In those last 20 years, computers became a formative accelerator to understanding all elements of intelligence and with the collected data increasing steadily and the computer power exploding, the changes will also necessarily accelerate.

One hundred years ago, nobody could have possibly predicted where we actually are today. Most of the central technologies powering this expansion of knowledge were not even invented yet. It should be obvious that in addition to the continuing and escalating prowess of human ingenuity, a completely new category of data processing will now be added.

Data collection is visible everywhere and is escalating in both speed and depth and compute ability with thousands of stacked GPUs is also exponentially increasing. Both of these categories are not just increasing but are explosively expanding, even considered separately. We are sitting today in a world where more money and resources are dedicated to this growth than has ever occurred in human history, and that growth will continue to escalate.

Some Dilemmas

The expansion leading to AI was not decided by us - it was prescribed by the Duopoly - Big Tech and Government, with an ever-increasing rate. Big Tech collects the data and provides the compute power, and Government makes the laws allowing this to happen. The Government benefits with money and data, assuring they stay in power.

Lots of discussion and fighting over how to control this runaway scenario will happen in small amounts, but the general trajectory will continue. The money and manipulative control available is unequalled throughout history as documented herein.

Data centers will continue to proliferate exponentially, collecting searchable records of virtually every action humans perform. All of this is new and is growing faster than any technology ever has grown throughout history, and the rewards for being the administrator of this tech are nothing less than control of humanity's direction.

As mentioned previously a huge dilemma is that, in truth, Big Tech is the only current structure capable of both collecting and using this data. If you had all this accumulated data in your basement it would do you no good. You can't use it in any way without Big Compute and the algorithms necessary to process it.

Likewise, if you had all the computers from your local data center in your basement it would also not be of any use without the data to operate on. Both are necessary to exist at the same place and at the same time and must also have access to all the electricity, water, and money to continue to operate.

It's a vexing dilemma we have arrived at.

As Einstein said in his Principle of Equivalence - if two things are indistinguishable from one another then they're actually the same thing.

Is this true also for the coming eventuality of a humanoid robot? The Van Gogh Principle says yes.

More on Exponential Growth

I have been mentioning the term “exponential growth” often throughout this document. It's really important to fully understand the implications of that growth pattern in order to relate it to the compute power explosion available to AI.

Let's use a checkerboard as a launch pad where there are 64 squares in an 8x 8 matrix. If I give you a penny on square one and give you another penny for each square you move, when you reach the last square you now have 64 cents. This is “arithmetic” growth and is a useful pattern to understand as it relates to many processes in life as we add experiences over the passing days and perhaps learn something new each day.

“Exponential growth” is another animal entirely. Start with that same penny on square one, only now I will double your pennies for each new square you visit. After eight squares, you now have \$1.28. Not bad and certainly better than the 64 cents you had on your last trip. As you reach the 16th square, you will now have \$327.68. We're just now getting the hint at the power of exponential growth.

By the 24th square you now have \$83,886.08. ---- Get ready for this – at the last square you

have over 92 quadrillion dollars - that's a 92 with 15 zeros behind it, and the jump from square 63 to 64 added over 46 quadrillion dollars to your payout!

This illustrates that not only does the overall number grow to unbelievable levels, but the growth amounts explode in size. On square one you have 1 cent and on square two you have 2 cents. By moving one square you have doubled your money. Your wealth has doubled, but you only added 1 cent. From square 63 to 64, you have again just doubled your money, but you just added 46 quadrillion dollars by one jump!

This is the real-world effect of exponential growth and this is the scenario taking place in the data centers as we speak. Thousands of GPU chips are being stacked into networks that again multiply their task handling abilities exponentially.

The GPU chips which ignited this data handling revolution were originally used to handle the task of speedily updating massive graphics images as necessary for your TV. In the data centers, they have been repurposed to “look at” vast amounts of data and locate any hidden patterns. There is currently a new generation of GPU-type chips designed specifically for this new application that, of course, will outperform the previous GPU generation.

Each of these steps along the way create exponential growth patterns on top of the previous exponential growth. Imagine that checkerboard on steroids.

Some Big Numbers

In my flight sim computer, my NVIDIA 5090 GPU chip can perform about 200 trillion calculations per second. Two years ago, the largest data centers had about 10,000 of these stacked GPUs. Meta, Microsoft, Amazon, Google, and Oracle are currently expanding into the 100,000 region with 600 billion already spent on chips in 2026.

To get a handle on the scale of this growth, and by extension better understand the exploding compute power, let's do a dive into some large numbers.

If you remember, “scientific notation” is used as an easier way to handle very large numbers. Essentially it is a number with an exponent denoting how many zeros following the number there are so, for example, 10^1 equals 10, a 1 with 1 following 0 - 10^2 is 100, a 1 with 2 following 0's - 10^3 equals 1,000 - 10^4 equals 10,000, etc.

Sooooo - 10^3 = thousands / 10^6 = millions / 10^9 = billions / 10^{12} = trillions / 10^{15} = quadrillions etc....

In our checkerboard example, we went from one penny in square 1 to 92×10^{15} dollars in square 64. For comparison, an entire human DNA string has about 2×10^{11} atoms - so my one GPU 5090 chip alone can do hundreds of computations per second on that entire DNA strand involving every single atom.

Of course, it's all far more complicated than just calculations per second but the illustration here

is just to get a real handle on the scale of compute power coming online and ultimately powering AI.

Just for a fun reference – The entire human body has about 10^{27} atoms- the entire earth has 10^{50} atoms-the Milky Way has 10^{68} atoms and the entire observable Universe has about 10^{80} atoms.....

As the GPU chips get multiplied and stacked into deeper layers and linked into ever greater networks, the computational power is quickly rising into these stratospheric numbers. This is a completely new realm of processing power and is still in its infancy.

This isn't just, “wow those are some big numbers” type of growth. These numbers have already outgrown the language we use to express numerical size. Billions, trillions, even quadrillions are being dwarfed as evidenced by the necessity to use scientific notation.

As mentioned previously - the human body has 10^{27} atoms. We don't have a word for that number because there was no need to express that number for any other reason than to denote that number of atoms. The computations capable of being done daily in a 100,000 GPU data center is approximately that same number! The number of atoms in the human body will remain the same - the compute power is soon to be able to exceed this and will continue to grow.

This grand scale must be continuously alive in our minds as we try to understand and predict where AI can go.

Coarse-graining

Allow me another little aside, as it may be beneficial.

A common practice in scientific reasoning is to employ a concept called “coarse-graining”. It essentially means that by looking at a system as a whole, you may be able to understand the operation of that system without the need to observe or even understand the individual, minute parts of that system.

Our Van Gogh effect idea employs this coarse-graining process. By looking at the copy of the Van Gogh painting, there is no need to look at each pixel individually, the important phenomenon is that once the copy is indistinguishable from the original, its effect on the world is the same as the original.

This coarse-graining technique of observation permeates all fields in science. Take, for example, a box full of air and insert a thermometer to measure its temperature. Whatever reading you get on your thermometer is essentially the average temperature of all the air molecules in the box. Some are warmer than your measure and some are cooler, but your coarse-grained reading is very useful in all thermodynamics computations relating to volume, temperature, and pressure. No need to have a readout of each molecule of air in the box in the same way that there's no necessity to define every pixel in the Van Gogh copy.

If you think about it, we live our lives by coarse-graining everything. Hidden just behind every single sight and sound we experience there are billions, trillions, or quadrillions of minute pieces

making up that event. For our daily lives, all this complexity is unnecessary in our coarse-grained view of the world.

So far, in human history, we have tried to make sense of reality by observing all these coarse-grained events as they occur. Scientists probing further into the depths of reality begin the trek into that complex underworld of minutiae.

The point here is that in all of human history there has been a hard limit on how deep you can travel - due mostly to incredible complexity and a limit on computational power.

Remember the checkerboard and the jump in your wealth from square 63 to square 64. Computationally, we are already on square 64 moving quickly to 65 and beyond. This growth translates directly into how deep you can probe. A lot of truths, so far unavailable to humanity, can be ferreted out of all those minute data points and is, as of yet, an unexplored field.

This is exactly how new progress will be made in all the complex fields making up reality. It may not be very important to be able to determine the temperature of every air molecule in that box or teach us much, were we to know that - but the ability to peruse the minutiae in a DNA strand is far more important and will likely yield untold new discoveries.

The incredible compute power we have been explaining here is beginning to approach the level of the task it is being assigned, the world at the atomic level seen as a vast data field.

Doing lightning fast parallel computations is one aspect of this process, but it's not the only puzzle piece. These computations need to be done “on” something and that something is data. Essentially everything is data - so how much data is there?

Data – How Much ?

There are about 8.1 billion people on Earth, and they cumulatively speak about 84 trillion words per day. An average book has about 100,000 words, so each day humans say the equivalent of 840 billion books. People have remarked that it would be a huge task for a computer to digest all the books ever written, when actually that is a small task compared to just monitoring words spoken daily.

As you can see from this example, the data currently being collected and saved is unbelievably massive and as we have noted, continuously increasing. Not only is the amount of data massive, but the compute power necessary to digest that data is even more so.

A lot of data is not time sensitive, so it can be stored and processed at a later time. Medical data, for example, can be collected as time goes on and viewed later to search for hidden patterns. Some data would benefit from quick or even immediate action. Military tactics would likely need swift analysis and current drone operations are already guided with this nearly immediate processing.

In the previous section I tried to illustrate with all the examples of huge numbers, the almost unimaginable growth in computer speed - you must now extend those numbers here to represent the amount of data gathered.

In actuality, the data even currently being gathered already far surpasses the ability of the compute power in all the data centers combined – but - even with that mismatch between data gathered and the compute power necessary to process that data, the current output from the AI models we can access is already quite amazing.

ChatGPT would easily fool Alan Turing into thinking that it was actually a real human he was interacting with. Suno is already composing music that is easily fooling almost everybody. Google can easily feed you data pertinent to your tastes by monitoring enough of your “output” to correctly make those judgments.

For practical daily use, the current processing abilities are enough to fool us completely by appearing conscious and smart.

We covered the idea of “coarse graining” because that is what we do when we think and the current AI models are also using that process to construct output to its users. As long as it can produce output that is indistinguishable from a real person, the illusion becomes reality. Since most human thought is based on these coarse-grained observations, it becomes a far easier task for the AI model to mimic this behavior. The data available to the AI model contains all the minutiae, but full assessment of that data is not necessary to model human intellectual skills.

As time goes on and the compute power continues to expand and the collected data also increases, the abilities to dive deep will also increase.

Course graining will continue, of course, but the ability to begin processing all the deep-dive minutiae will escalate as the exponential increase of computer power continues. Some of the data that will continue to be collected will be repetitive, and this fact together with the course greening will allow the steadily advancing computer power to come closer to keeping pace with the data collection. Of course, the closer it gets, the deeper it can go.

It's important to note that data collection, although massive, grows “additively” while compute power grows “exponentially”.

Humans are fragile creatures with limited lifetimes and even more limited productive years. Most humans also don't care much about minutiae; even in history, only a few even considered it - so few, in fact, that we know their names: Aristotle, Newton, Einstein...

Since throughout all of history the most successful people were the best coarse-grainers, it's not unusual that focusing on the minutiae of life was not an important evolutionary development.

Aristotle thought that all of reality was composed of four elements: earth, water, fire, and air. This was the best coarse-grained view at the time, about 2,300 years ago. As science became able to peruse the minutiae of existence, “atoms” and “elements” were discovered and replaced Aristotle's ideas.

This illustrates two salient points. One, as more minutiae is discovered our understanding improves and, two, coarse-grained conclusions, even ideas held for many years can easily be wrong.

The diversion into Einstein's relativity contained herein is an example of this. For all of human history, we have considered space and time to be part of some unchanging underlayment to life and the events that create history. Space is the same everywhere and time passes at the same rate no matter where you are.

This is the coarse-grained view that is held even today. All our daily experiences support this coarse-grained view, but of course it is wrong. In our daily lives, it is unimportant that it's wrong, mainly because it's unnecessary to go deeper into the minutiae that show exactly how it is wrong. As long as our GPS receivers work, we don't have to consider that the inventors of that technology had to probe that minutiae, understand it and assemble some electronics that can be exact enough to locate our position precisely.

This illustrates another important point. Our coarse-grained conclusions continue to evolve as the collection and understanding of minutiae increases. What we know from today's coarse-grained conclusions is more accurate than Aristotle could do with his more limited access to minutiae.

Even though it is often unimportant for daily life if we understand the minutiae underlying these advancements in human understanding, it is of course of prime importance that someone do the probing. Advancements in all fields of science rely on understanding this underlying complexity.

As we know more and can probe deeper, the depth of this complexity grows really fast, maybe exponentially. With this prime reliance on correct assessment of the complex minutiae underlying every part of reality, does this not call into play the AI technology that is the most capable method to peruse this vast playing field ever to exist?

A Quick Re-hash

Let's do a quick revisit on some of the ideas that we need to deeply understand about what AI is now doing and where it can go.

Our observations on the progression of human awareness are important. We mentioned Aristotle and showed how a subsequent deeper understanding of the minutiae underlying his ideas exposed the errors in his coarse-grained views. We also mentioned Newton, and the same progression happened with Einstein's discoveries. Its important to remember that the quality of our coarse-grained views improves as the understanding of the complexity of the minutiae underlying it increases.

We have covered a lot of territory and hopefully have some tools at hand to try and craft an informed understanding of AI, at least as far as we can given the minutiae that we can employ in our coarse-grained judgments.

Even though our coarse-grained judgments can be improving, it should also be remembered that coarse-graining always contains errors due to the limited understanding of the complex data underlying it. We are currently also dealing with the fact that much of the underlying complex data is hidden from us and in the hands of the Duopoly of Government and Big Tech. We don't really know what and how much data is being accumulated and only generally can understand the compute power being

assembled.

As stated previously, it's also likely that the complexity of this whole AI structure is beyond any of us to really understand. Remember Demis Hassabis' statement 10 years ago when his AlphaGo/Master program beat the reigning world champ, “the computer is making moves no human has ever made and that we did not program and therefore can't predict”.

Ps: Hassabis is currently the head of Google's DeepMind division.

Let's push ahead anyway while we also keep this in mind. Where is AI today? Let's extend that re-hash to what we already know.

We have computer technology steadily expanding by huge amounts. We tried to understand the scale of this expansion by probing some huge numbers, from the checkerboard growth of our pennies to denoting the number of atoms in structures up to the size of the entire universe.

We also have the same advancing level of data being collected and stored. Virtually everything we do is or can be monitored, with the devices and methods of collection expanding into all aspects of life.

At this point, the only entities that can collect and store and process this data are the Tech companies. Their reach is essentially into every aspect of life – internet, media, cameras, smart everythings, texts, email, phones, satellite monitors, etc. Essentially, this collection is unmonitored and uncontrolled with no meaningful limits on its collection or usage.

There are some laws governing some aspects of privacy, but does anyone believe they are being obeyed? As we have also previously discussed, the value of that collected and processed data to both Government and the Big Tech companies is huge.

The Tech companies make money monetizing every aspect of that data and feed some of that money to Government, who not only enjoys the monetary rewards but can also use the processed data to manipulate the electorate in many ways we have explained previously. Once you can manipulate the “inputs” to people and monitor the “outputs” of those people, you have a foolproof feedback loop control network. The current dilemma is that the only group that can do that is the Government-Big Tech Duopoly.

I do not see how that Duopoly gets dethroned. The Government gets the votes and Big Tech gets the money by monetizing the collected data. Big Tech manipulates the info available to us both to enrich themselves but also to swing public opinion so the current Government can remain in power.

The Government responds by making infrastructure favor Big Tech expansion and not addressing the obvious problem of intrusion of privacy. The more data collected, the more both factions of the Duopoly benefit.

As stated previously, there are essentially no current limits on data collection and no meaningful public discourse on any topic that could limit it. Even where some laws may exist, nobody can police compliance since the only capable entity is Big Tech, which is one of its benefactors.

It's the ultimate case of the foxes guarding the henhouse.

Even if we take all the data, it does us no good if we can't process it, and if we take all of the compute, we would still not know how to use it. We know it's not a good idea to let any one entity have all this unmonitored power, but they may be the only ones that can actually run that complex set of technologies. That's a major dilemma.

We can also ponder whether we humans should even be pursuing the AI grail. I believe we will continue unabated for a couple of reasons. The first is that it's not only the most profitable technology ever devised as far as actual money is concerned, but it is even more valuable as a social control mechanism.

As far as controlling AI goes - it is a fool's errand to let people who do not understand the technology try and control it. They would not know what to do or even know if what they are doing is effective. It's also the case that technology this powerful will likely always be kept secret, even now, the most advanced tech is already secret. No reason to expect that to change.

Since a big component of this AI field is "Government", we are forced to peruse the political domain. Currently with 40% of the electorate unwilling to communicate with another 40%, we enter a minefield dominated by anger and self-righteousness – with virtually no quality discussions between the factions.

Without meaningful intelligent problem solving, there are no disagreements that can be mitigated. Add to this the facts as presented herein, that the growth of AI is massively complex even to those directly working in the field.

This is not a good precondition to foster intelligent conclusions.

One entity, the Duopoly of Government and Big Tech, currently controls all the collected data and are the only ones with the compute power to manipulate that data. They also own or technologically empower nearly all the means of information dissemination. These are all the pieces of the puzzle necessary to exert the social control we have been discussing.

Since we are trying to understand where we are along this AI development timeline, we should be looking around us for clues. We don't have access to the collected data and we don't have the compute power even if we had the data. We don't have access to the algorithms used by any of the search engines that currently decide what we are allowed to see. We also don't know what information has already been removed from our view by Internet providers.

All of these procedures and more are silent, hidden components necessary for group thought control. We think it's okay that Google feeds us Home Depot pop-ups when we're looking for patio lights or maybe cheap Delta flights to the Caribbean when we are thinking of vacations.

Let's remember that Google sends you those specific ads because they collected your data and knew what you were thinking.

The data they collected and processed is what allowed them to make a profit from your future actions. The more data collected and the better their processing algorithms, the more profit they can make by motivating your actions.

Since Google already has the mechanics necessary to make this information control work system-wide, it's logical for them to see where else this may work to generate even more profit.

In truth, the money Google can make from Home Depot or Delta is minuscule compared to the money available from Government by implementing social control. It would be a bad business decision for them not to expand into this lucrative avenue !

For example, if you can keep wars happening, here's where the money goes - Raytheon (RTX) \$58 billion - Lockheed Martin \$70 billion - Northrop Grumman \$40 billion - Boeing \$89 billion.

Controlling public sentiment here alone pays way better than Home Depot possibly can.

The defense industry alone also contributes over 100 million to selected political candidates and PACs during each election cycle. A little nudge from targeted “inputs” can easily swing huge amounts.

It's also clear that the companies that contribute these massive amounts to political groups expect a return on investment, and it's obvious that they get it since the contributions continue year after year.

The Duopoly sits at the exact center of this scenario.

Government gets the huge contributions, and Big Tech, using their collected data, tweaks public opinion by selectively adjusting data to the electorate - and in return, they are assured their monopoly in data collection and favored status, allowing their unobstructed expansion.

In the U.S. now, there are many large organized groups that funnel huge amounts to various political entities. No national or even statewide candidate can attain and maintain power without massive cash inputs by the constituents that stand to gain by their actions.

Money in politics has always been the case, but AI is perfecting the pay-for-play process. The effectiveness of “tweaks” to public opinion are now easily measurable with real-time data collection. Adjusting the “input” to get the desired “output” is getting easier, faster, and more efficient.

We are approaching “Techno-Authoritarianism”, the use of digital tools by Government or powerful elites to monitor, manipulate, and control populations.

At this moment in time, what I have called the “Duopoly” is Government and Big Tech in somewhat of a mutually beneficial power balance. Government gets money and stays in power, and Big Tech maintains their monopoly and monetizes their data collection. This Duopoly lives through all administration changes and continues its expansion.

Big Tech maintains their favored status by providing the processed data to Government to implement better and better social control. These are still the early days of the escalation of AI

abilities, but as they grow and mature, I think the Big Tech group will soon overpower the Government sector.

The collection of data and the manipulation of that data all lie in the Big Tech domain. At this point in time, they still need Government to allow and expedite their explosive expansion but I believe that at a certain point, Big Tech will realize that Government needs them far more than they need Government.

Big Tech has the means for the end game (societal control) and can manipulate and adjust public opinion on their own. As this moment approaches having favored status from Government is no longer necessary - Government is now the junior partner.

That's likely a wildly different future than anyone has been expecting.

Getting back to our central theme of AI and realizing that since we have Government and Big Tech joined at the hip, where are we now and what may happen next?

As stated previously, politically 80% of the population is divided into two equal-sized and vehemently opposed groups. Each 40% is so sure that their ideas are correct that they will not even talk to the other 40%. The remaining 20% is a diverse group and has no collective power and no large amounts of targeted money to offer compared to the larger 40% groups.

The issues underlying those vehement disagreements are, in reality, often so complicated that few in either group can explain in much valid detail how they arrived at their current conclusions except by invoking their “feelings” and further with no meaningful discussions occurring there is no possibility of a mutually agreed upon resolve to any dilemma.

Of course, in our “free” society everyone is entitled to their opinions and allowed to speak or write them, but let's not pretend that unthought-out conclusions based on inadequate evidence have much utility to solve any real problems. “Social consensus” may help alleviate the mounting vocal disagreements and function well as a peacemaker, but will not advance the goal of underlying problem solving.

Problem solving discussions should be taking place mainly between parties that disagree - but it's easy to see that this is not happening thus removing the only mechanism for the mitigation of differing opinions.

Take any current topic where there is disagreement between the 40% groups – Immigration, for example. We have seen that when one of the 40% groups is in power they let huge amounts of unvetted people cross the borders – and when the other 40% is in power they try to remove these same people. Neither of those policies is well designed or implemented.

One 40% side believes that since we are a wealthy country and care about the well being and safety of others, we should be willing to help those whose situations are not as comfortable as ours and let anyone come here who wants to. They also likely believe that the free mixing of cultures is a net benefit. The other 40% side thinks we should control the borders and restrict who can come here

because too much uncontrolled immigration decreases the quality of life that we have spent many years creating and that mixing high-trust societies with low-trust societies creates more problems than benefits. Both these positions are legitimate concerns.

Of course the only way to decide on “one” workable policy is for those 2 groups to talk to each other and somehow agree on an effective compromise solution. This, of course is not happening so we careen back and forth every new election cycle – neither policy working,

Looking at human history with a critical eye, this 40-40 split may be more the rule than the exception.

Who benefits from this 40 – 40 split. Certainly both of the individual 40% groups do – collectively they are the Government. Each 40% creates their own “in group” of like minded people and can join together and “fund raise” to beat that other hated group – but who else benefits ?

The Big Tech companies do. Not only are they in at least partial control of the data that each group can access and therefore able to manipulate opinions - but the continuous anger between the sides keeps anyone from looking at the bigger issues. All the energy spent on fighting the other group is energy not spent on looking at the more important (and more complex) issues. The Duopoly is able to continue its expansion largely unobserved and unchallenged. They have every interest in keeping the 40 – 40 split energized, uncommunicative and angry.

We are in the depths of “politics”, and anger and distrust become the operant motivators.

As in any “problem resolution” scenario, “clear and honest” discussion is a necessary starting point and this requires full participation from both sides especially since the dilemmas we face at this moment in history are massive and complex.

One thousand years ago, the most damage one crazy person could cause was to hit his neighbor with a club. Two hundred years ago that same person had a gun and could now hurt anybody in his village. Today, there are buttons that one person can push that could annihilate most humans on earth.

The stakes of these “clear and honest” discussions are considerably larger than at any time in history. Given what we have discussed in this document, what can be done to fix this dilemma?

First, we should realize that just because we want to solve all these ongoing problems, there may not actually be a solution.

At least part of mankind has always been on a quest to live peacefully with the rest as far as we can see back in history yet disagreements leading to violence and wars have been omnipresent and continuously recurring.

In order for lasting peace, everyone has to be on the same page, and that seems at least unlikely and may be functionally impossible.

Attacking this dilemma with our most tactical approach, we must realize that at this moment in history, we must take the “complexity” issue very seriously.

If you need your refrigerator repaired, you must find someone that understands how that refrigerator works. You want someone to design a safe nuclear power plant - you look to find the best trained nuclear scientist to do the job. The reason we look for people trained in these particular disciplines is that the “complexity” of the task is beyond our personal understanding.

An efficient elected Government should be working with this in mind. The people we elect don't need to fix refrigerators, but they do need to operate in a vastly more complex matrix, encompassing basically everything important to ongoing life, now including making informed judgments on AI.

It's difficult enough to find a good refrigerator repairman - good luck finding someone competent to navigate the diverse and complicated environment of today's life and the choices that will affect tomorrow.

The brutal truth here is that there is no refrigerator repairman for AI !

Truth is “expensive, difficult and often painful”. Expensive because it takes huge amounts of energy to dive deep into complexity. Difficult because that deep dive trek is not easily defined, has no defined endpoint, and often the best conclusions you can find are not very fulfilling.

Fiction, on the other hand, is easy, cheap, and often emotionally pleasing. This is the reason that “Bread and Circuses” is so easily employed and has been an effective procedure used to placate whole societies. Unfortunately, it is the opposite of what should be done to persevere in truth-seeking.

I'm afraid it's probably unrealistic to expect a significant number of people in any society to come up with the large amount of energy needed to attack the increasingly complex issues currently on our plates and, as the underlying technology increases in complexity, it becomes even less likely that any of us can keep up effectively.

Here's a partial list of how deep data gathering has penetrated into people's lives.

Water systems / electric grid / well pumps / vacuum cleaners / ring cameras / smart everythings / phone systems / calculators / satellites / GPS / cars / gas pumps / Visa-MasterCard / flashlights that blink / blenders / digital clocks / everything wireless / wall LED dimmers / coffee makers / all USB connections / and children's blinky shoes. Anything with a chip already is or can be networked or internet monitored. This is the “Internet of Things”

The easiest undeniable proof of this advanced data gathering is to just watch anyone under 40, on the street, in a store, or even at home - they are all glued to their smartphones during every waking moment. Some of this can be attributed to normal curiosity, but the obsessive addiction has obviously been planned to make you stay on that virtual location as long as possible. Of course when asked publically if this is part of their plan the Tech groups all deny it and are let off the hook by their main partner – Government.

If you're on TikTok or Facebook, the longer you stay, the more likely you will see that Home Depot ad and the more they can charge Home Depot. That is certainly a measurable factor, but the real

value currently is the data that they are gathering.

If you're on an iPhone, Apple knows everything you are doing. Not just which sites you're visiting, but your location and what you might be saying verbally. At Facebook, Meta also knows not only what you are actively looking at, but knows all your friends and contacts and what they are looking at. Meta has also huge amounts of data on what your relationships are with all those friends. On and On and On....

With AI now maturing, all this is being collected and stored, sometimes for immediate use and sometimes for future reference. Since computers don't need to sleep that "future reference" can easily mean "forever"

So where does AI enter into this picture?

As I have tried to explain in this document, AI is already firmly and deeply entrenched in the world of today. It is progressively gathering more data and adding the compute power to process that data faster and probe deeper into the levels of the minutiae.

It's been interesting to ride the spaceship with Andy and watch people eat grapes. Each of these thought experiments, hopefully gives us new data from which to extract our best coarse-grained opinions.

On the political side, in the Western world we need to find ways to deal with the Duopoly behemoth or become a more and more programmed society with applied AI using all the manipulation techniques illustrated herein. We are clearly already on this slippery slope and seemingly accelerating in our slide.

With the deep penetration that has already occurred, I'm not sure how to stop this but I do know a couple of necessary skills we must employ to even have a chance at fashioning a better future.

Number one is we cannot be intimidated by complexity. It is and will be a growing component in every important topic necessary for valuable problem solving. The expanding checkerboard lies ahead of us and will be present whether we look at it or not.

Number two is that we must engage in all types of serious discussions, including political. At least in the current world, the political component is the only arena of public problem-solving. Without "going there", the careening from one polarity to another will continue unabated.

Ask yourself, "when was the last time I had a good discussion with someone across the political divide?" The answer is usually "never", or at least "very seldom", and in fact most people actively stay away from this arena.

Almost all political discussion is with people you agree with. This is far easier and more pleasing, so it's not surprising, but unfortunately does nothing to enhance problem-solving. Actually, most beneficial discussions should be with those you do not agree with, but they are "expensive, difficult, and often painful". That's why we don't do it.

In the end, there may not be viable agreements available, but without the deepest discussions that we are capable of doing, there is absolutely no chance of resolutions. You can still let your feelings guide your opinions, but without open discussion with other humans, no mutual ground can possibly be discovered.

Truth is expensive, difficult, and often painful - but it is the only road that has any chance of real success.

Synopsis

For our last run at making sense of all we have covered, let's do a quick synopsis of where we have been. I have tried to include here a wide range of elements that I think are relevant to making an educated guess at what AI is, what it can do, and where it can go.

----- First we looked at computers in music, starting with some history of technological improvements over the years that shaped the evolution from Beethoven to today. I used “music” as an introductory topic because it is an “art”, and so far has only been done by humans. I tried to show that “magic” is easy to recognize if you have access to all the “outputs” of the listener, and that can easily be done by a computer.

I defined “outputs” as the sum total of whatever a human does when presented with a stimulus, and further defined “inputs” as whatever that stimulus was. This “input” and “output” terminology permeated this entire document as it is a good way to understand what AI may use as one of its operational data handling procedures.

----- This led to the idea of the “Van Gogh effect”. The idea is that for AI to do any task, it does not need to be perfect; only good enough so we can't tell the difference from whatever we are comparing it to. This idea also shows up regularly throughout this document.

----- Another idea that shows up continuously throughout this document is the concept of “coarse-graining” and the realization that we live our lives this way....We must simplify the barrage of minutiae that is omnipresent to our thoughts and senses in order to reasonably make our daily decisions. We must also remember that “coarse-grained” decisions are often wrong and require that we update them as we learn more of the underlying facts.

----- I delved into a bit of Einstein's discoveries and even some Quantum Mechanics to further illustrate how the “coarse-grained” conclusions of even Newton were wrong and were updated with new discoveries. This also illustrated that in some cases one new fact can sometimes change entire trajectories and I postulated that maybe AI is one of these new facts.

----- Starting with the section on music and continuing throughout, I tried to explain a bit of computer technology and its rapid exponential expansion of abilities. I explained how GPU chips empowered this expansion into new levels that were simply not possible even a few years ago. In addition, neural networks and parallel processing opened doors to handling incredible amounts of data.

----- This led us into a section where I tried to get us comfortable with some BIG numbers. Not

just BIG, but gargantuan, explosively BIG — size of the universe measured in atoms, BIG. We don't normally encounter these numbers in daily life, but the computer explosion in power and speed is moving into this territory where we need these numbers to explain how fast processes are occurring even now in the functions of the computers powering AI. The “checkerboard” example was included to try and illustrate how exponential expansion can easily grow to a place where these BIG numbers are also necessary to explain the results.

----- This led us to a discussion of the AlphaGo saga, again needing these BIG numbers to try to understand how a computer can beat the best human player on earth using moves no human had ever done before. Another conceptual door was opened when the Master program beat the AlphaGo program, since Master was entirely self-programmed. This “self-programmed” method is emerging as a wild card in any realistic discussion of AI capabilities, especially as time moves forward.

----- I introduced you to “Andy the Android” and let him eat some grapes. We were trying to make some judgments on what is necessary for him to fool us into thinking he was a real human. I tried to further the idea that the only way we can tell if he (Andy) liked the grape is to watch his reactions, observe what he says, and note the expressions he makes etc. Since a computer can easily collect that data, and if that computer was housed in a realistic human form like Andy, we would have no way to tell it was not a real human. Andy fooled us - it was the “Van Gogh Effect”, only with a copy of a human instead of a painting. We discussed how this is also related to how AI could make “magic” in music - adjust “inputs” until the monitored “outputs” exhibit appreciation.

----- All these areas we visited relied on not only expanding computer technology but also data gathering and storage procedures. The penetration of data-gathering sensors and monitors in everything down to blinky shoes was discussed, hopefully culminating in a fuller understanding of the huge databases being assembled. This led us to understand that data centers are the repositories of all this data and also house the computers necessary to make use of that data.

----- I tried to show that currently, the only entity that has the money, technology, and computing power to collect, store, and process this data is Big Tech and further, they need the Government to allow and help create the huge infrastructure mechanisms necessary for exponential growth to continue, and to allow the ever-increasing collection of personal data to expand.

----- This creates what I called the “Duopoly”- the merging of two monopolies: Big Tech, held in the hands of only a few multi-billionaires, and Government, which partakes in both the money generated but also increasingly the societal manipulation that AI can create. I contended this is a big problem without an obvious solution. Each party benefits from the other, and both have vested interests in keeping their operations private.

----- We also had occasion to bring “politics” into this discussion because Government is one half of the Duopoly. This also forces us into a discussion about how “choices” can be influenced by the abilities of Big Tech to create “inputs” and monitor “outputs” of individual humans. This was the method AI used in music to decide if a song was as good as McCartney's. "It's vitally important to decide where this is now, where it can go, and where it should go."

----- Finally, and I believe most importantly, we tried to understand “complexity” and how it relates to AI. We humans “coarse-grain” virtually all our thoughts because it is functionally impossible

to include all the minutiae of any moment into our conclusions, even as we must understand that along with “coarse-graining” comes the almost certain inclusion of errors. If our goal is a smooth, pleasurable glide through the day, then it's “coarse-graining to the rescue”. If however, our goal is to really understand a topic, we are forced to go as deeply as we can into the complex jungle lying beneath everything, all the while realizing that no matter how deep we go, we're still coarse-graining the level just beyond. Understanding the size and complexity of that underlying jungle is why we had to delve into the “huge numbers” section we discussed previously.

All the individual parts I have presented here are pieces of that complex minutiae that are necessary to help understand AI. Some are hardware parts and some are conceptual parts. Some are self-evident, and some are just my opinions and added to that should be your observations. Taken all together, they are still only a minor subset of a huge number of relevant points we should be including in any trustworthy conclusions about AI.

I realize there is a major dichotomy here in my presentation. In one section I contend that we must have discussions focused on all our disagreements as this is the only method for resolution - and in the next section I say that everything is either too complex to understand or that much of the pertinent data is either unavailable or pre-tuned.

Unfortunately, both are correct and lead to a dead-end of sorts. This is the perfect example of the fact that real deep thought is expensive, difficult, and often leads to troubling conclusions.

It is expensive for our brains to collect these troubling premises - it is difficult to see any easy fixes to these dilemmas, and taken together, they lead to either no conclusions or at least troubling ones.

One of those troubling conclusions must be clearly stated - just because we can frame the debate by illuminating all these disparate premises does not mean a conclusion is possible. There is not a clear logical route through the aforementioned points that can resolve the dilemmas.

All of the above elements in one form or another have been a part of life for a long time and the interplay of powerful factions is nothing new in human history. The only new agent embedded and underlying all the above observations is advancing AI capabilities.

We are currently on the road to Techno-Feudalism unless some unforeseen events redirect this momentum. Perhaps this is where the current conversation should center.

With AI increasingly embedded deeper and deeper into every aspect of Earth's systems and soon to be the only possible entity that has any hope of comprehending the increasing complexity, perhaps this “New Director” is inevitable and we are just along for the ride.

AI Going Forward

Sooooooo... let's see where our coarse-grained guesses at deep truths lead us with all this data on our plates and also let's realize that AI will change quickly over time. I would like, therefore, to separate current times from future times since we are actually here to see and experience these current

events at least as far as we are allowed to see, and, of course, only to the limits of our perusal of the infinite minutiae underlying all of reality. We will then move to the conjectural side of things with the AI future separated into short, medium, long, and then maybe really long-term timeframes.

Current

With “complexity” being a necessary and omnipresent component in all these observations and conjectures, we must remember the mantra: “it's expensive, difficult, and may lead to troubling conclusions”.

The playing field of AI is nothing less than the sum total of all human experiences arranged into a searchable database, including all “inputs” and “outputs”. The “outputs” are all observed and saved and the “inputs”, adjustable by the AI models.

Currently, this entire structure is privately owned by Big Tech and is unlimited in its pursuit of data collection. It's also obvious that our Government allows and provides the necessary infrastructure components to be put in place, again, with no meaningful oversight or public debate. Big Tech gets this preferred position by providing the Government with both money and data that allows for social control through Big Tech's control of “inputs” and “outputs”.

In this discussion, I have tried to illustrate that the incredible computer power paired with the mammoth data collection already occurring gives the AI models a larger playing field than has ever existed before. The amount of minutiae that is already stored and can now be probed is unprecedented at any time in human history.

All these pieces are here today to some degree. Certainly, computer power is escalating, and massive data collection is evident and both of these are necessary components of why data centers are expanding. The huge amounts of money being spent there are a visible testament to their central importance to AI and the Duopoly.

The social control aspect of AI is still a question as to how far it has progressed. I think it's undeniable that Big Tech and Government both have visibly controlled the availability of some data. It's for each of us, unafraid of complexity, willing to do the difficult analysis, and unafraid of troubling conclusions, to make our own judgments on how far this manipulation has gone.

Wherever your judgment takes you in that arena, we still arrive at the current fact that Big Tech and Government are the sole benefactors, both in the money generated and the control available.

It's also pretty obvious that Government “corruption” is endemic. Northrop Grumman and Raytheon would be out of business if wars were to stop. They sell bombs, and unless those bombs are used, they don't need to be replaced. It's no surprise these businesses spend huge amounts of money supporting and lobbying legislators who favor the ongoing wars and once those legislators retire, they often are given plush, high-paying jobs in the very businesses that profited by their previous decisions.

Other big business conglomerates also contribute huge amounts to the same Government related individuals and entities – Big Pharma, Big Ag, Realtor Groups, Big Banks and Finance, Utilities, Oil

and Gas Companies, Insurance etc. All expect and receive some sort of “Return on Investment” or they would not continue doing so.

AI enters the metric because of its growing abilities in thought guidance and shaping through data manipulation. It's already being done in music, and the defense industry pays way better.

The centralized Duopoly control of all this technology is the current “elephant in the room”, or the “rat in the punch bowl” - whichever image you like better.

It's here now and growing by the day with no brakes - and the current political climate is doing the 40 versus 40 dance and the electorate is unable or unwilling to focus on deeper, more complex issues. How to change this - I don't know. You can't “vote yourself out” of these problems because the AI / Government Duopoly continues unabated regardless of which party is in power.

There is another AI technology that is already present and waiting in the weeds - that is the merger of computers and humans in a physical sense, either “chips in people” or “human tissue in computers”. Chips embedded in humans have already begun with pacemakers, cochlear implants, insulin injectors, artificial hearts etc. - but also recently with chips placed directly into the brain and while this is still in its infancy, it will likely progress quickly.

The reverse has also begun with actual human DNA being studied as a data storage device. I covered a bit of DNA architecture here and noted the incredible amount of information encoded there. That makes it fertile ground to explore since data is the bedrock of all AI's abilities.

This is where we are today. AI is an incredibly complex, promising, and yet terrifying technology that's likely to be unruly at some point.

Short Term

The short term of course starts with what we just outlined for today and just adds weeks, months and maybe even years. For my near future prognostications, I am NOT including any new technologies but instead only linear advancement on what is already here.

We see advancing data gathering and storage and compute power able to probe ever deeper into all available data. We have ChatGPT and other Large Language Models already able to act intelligently and take part in interactions that were never programmed by any human. Suno is already crafting music that is indiscernible from what a human might create. Early versions of “Andy” already exist that may soon pass the Van Gogh test and become indistinguishable from a real person. Even linear projections into the future from these various elements lead to stunning possibilities.

The fields of sociology and psychology can certainly be used to try and understand our human reactions to this changing dynamic, but these fields are not applicable to “Andy” or any other component of AI – they respond only to programming, not feelings or intuition.

AI has already permeated every aspect of human life and extrapolating forward in time, will continue both operationally and likely organically, as stated previously. Just look at Earth's children to

see how deeply AI has been embedded.

Children are the human component predicating the future and AI must be seen as a major formative element of any futures we may imagine - and given that human evolution is way too slow to keep pace with these events, AI may be the major element.

All of the examples I presented herein to illustrate how one new discovery changed everything that followed, and this may be the applicable model here.

On the good side, you have a mechanism to begin probing deeply into everything, and are increasingly able to handle the growing complexity of doing so. Problem solving in all arenas is enhanced and will certainly yield some useful new discoveries, including human-computer interfaces, including implants, and continuing research into DNA complexity.

As long as the Big Tech-Government Duopoly is fully in the driver's seat, you will see all or most of the discoveries tuned to their benefit first and then spread out to the rest of us if it's beneficial to them. This is the reason that any singular controller is not a workable long-term arrangement.

At this time and into the near future, the Government will not be a trustworthy partner in any possible removal of that power since they gain so much by keeping it intact.

A few things may change this metric. Certainly a monetary collapse may radically change the political structure, along with some revolution or mass rebellion. In the near term, I can't see revolution-type changes taking place because, in order to organize any action of that degree, you would need media and free immediate communication, both of which are increasingly controlled and monitored by both parties in the Duopoly.

Since we're talking about the future, we must also consider the possibility of apocalyptic-type events that would necessarily change these trajectories. Volcanoes, asteroid strikes, nuclear war, coronal mass ejections, or a visit by aliens, - any one of which could easily disrupt almost every structural part of life and would certainly drastically alter what follows. It's also possible - and perhaps even more likely - that less than apocalyptic events would disrupt or slow the expansion process: electric grid failures, rare earth mineral availability, or GPU chip manufacturing disruptions - all of which would cause at least a minimum disruption to the timeline.

Short of that, I think that the near-term future will vacillate back and forth as it is now – as all the AI capabilities continue to expand exponentially and the data centers continue to proliferate.

One more eventful change that may occur in the near future is the progression of AI “reach”.

So far AI has mostly been “used” by humans as a “tool”. In the near future given that the complexity of structures underpinning all of society will likely continue to increase, AI may increasingly be acting on its own as “Agent” - making choices and allocating actions that may not have been specifically programmed. In simpler times humans would have been able to handle these unforeseen events but as complexity grows it's easy to see how ultra fast computers would be the only control mechanism that could respond quickly enough. This is a type of “Breakout” but it is not the only type that can happen.

Here's another - The entire world is already interlinked with a web of untraceable connections. We almost daily hear about hacking attempts from locations all around the globe, proof positive that this web exists and is not entirely controllable. This is the “backbone” allowing the “breakout” risks to occur and spread to unpredictable places. .

Almost every infrastructure of any size is already managed and controlled by computers: water systems, electric grids, hospitals, stoplights, even your house has a smart meter reporting to the local system, and certainly your local data center has computer controllers allocating its own local power distribution, cooling controls, etc. Each of these systems was programmed to maintain optimal operations, hopefully including all possible unforeseen events that may occur.

If the hospital system is confronted with some power disruptions, the controlling computer would likely start its backup generators. If the power demand was more than they could provide, its programming would shut down the parking lot lights before the heart-lung machines in its critical care units.

Likewise, the local AI data center would also have its computer control network programmed with a hierarchy of importance. Confronted with a power outage, it certainly will also have backup generators and will likewise shut down parking lot lights before its main storage and processing units, which would necessarily take priority over everything else.

With this huge worldwide interconnected web we just mentioned, the difficulty comes when these systems become linked in some scenario that neither system was programmed to handle.

The hospital and data center computers each have their own hierarchical rankings within their own systems but not between systems. The original programmers of each individual system never assumed that the hierarchy problem between separated systems would become an issue since there was no interconnection at that time, but as we have seen these connections can sometimes occur without previous planning.

If both the hospital and data center electrical grids get stressed by some unforeseen incident, which one gets shut down first? And if AI is controlling both systems, it will likely place its own operations at the top of the priority list. Since virtually everything is interconnected, these unpredictable clashes will become more and more likely to occur.

These “breakouts” were not pre-programmed and can occur at unplanned moments and with unforeseen consequences. Once again the complexity of the myriad of interconnections make these events almost inevitable – even with the most careful of planning.

Mid Term

In guessing at the mid-term future, this is where the speculations really begin to expand. Humans are always changing, and “shit happens” when we consider everything that is possible. Big powerful events that can alter the status quo become more likely as time goes on.

“War” is currently employing all that AI can offer and will likely escalate to being the main component. The speed and depth of essential computations provided by AI will become the new description of conflict, not just in weaponry but also in cyber attacks, biological modifications, and even psychological arenas. AI is not subject to “biological” damage, so it may re-evaluate what is important for its survival, and psychological manipulation may well be one of its tools. Again, so far this is just a linear progression from what we can already see in our current environment.

The self-programming of AI computers that is maturing today will soon create “non-linear” unthought-of methods and even goals as we move further into the future. At any moment, there may be new chip designs that use far less electricity or run much cooler, easing the strain on infrastructure. Quantum computers could be developed escalating the reach of computation into vast unpredictable arenas.

Since AI's ability to self-program will continue to advance, and since we are not controlling either the speed or direction of that advance, it is entirely possible that AI itself could design these new computers or chips. Remember Demis Hassabis' words from over ten years ago: "The computer was making moves no human had ever done."

It is certainly possible that the Duopoly may be broken up or reorganized with these new AI abilities – OR - the Duopoly may continue to escalate the increasing control it has over us and tune us into easier compliance. Both “1984” and “Brave New World” contemplated some version of this possible future, and AI would be the wizard doing the dirty work.

As time goes on and AI abilities expand the previously mentioned “Breakouts” could also occur at any moment with unpredictable effects. Computers are now and will be embedded into every part of life and infrastructure, including human bodies. It's not hard to see how a network this large and complex may behave in ways we did not plan or intend. With self-programming increasing this “breakout” may even be unnoticed. Some companies have departments trying to control this safety issue but it may just be – “Computer vs. Humans” - trying to operate in undefined arenas. I know where my bet would be....

The unavoidable complexity of possible events over time makes anything other than wild guesses untrustworthy as to what the mid-term future could be like and, unlike any time in the history of humanity the variety of events that can shape this trajectory now includes AI as the ultimate wildcard designer.

Long Term

All the uncertainties in the midterm diagnosis will of course, continue into the longer term. If we are able to speculate about this more distant future at all, we must assume that AI development continues. This is certainly not assured, but if any of the apocalyptic events we have covered actually occurred, then society will be either at some nasty standstill or even gone, eliminating any reason to even be guessing at these distant futures.

Soooooooo - If we sidestep all the destructive scenarios and somehow muddle through the

midterm future, what is possible in the longer term?

Even with today's AI abilities, technologies that used to take decades now take days and the problems broached can be exponentially harder with increasingly more data being compiled. The ability of AI to self-program will certainly continue to widen the possible futures to ever greater degrees and data collection will have penetrated every aspect of life.

AI will see humans as huge data production machines and with many years of data gathering and processing already done on these humans, it's likely that many phenomena of life will be much better understood. Certainly, all the chemical processes involving life will have been subject to the massive pattern recognition abilities of future GPU chips. On the good side this could allow many diseases to be understood and perhaps cured. It's even possible that things like "aging" will be better understood and perhaps slowed or at least mitigated in some way.

It's almost assured that by the long term contemplated here, computer chips will be merged with humans in some yet undefined way. Computers are way faster than humans at processing data and have nearly infinite access to that data, so the interface would advance human abilities to untold levels.

AI may see that its addition to humans is an advantage. Humans may still possess some abilities or aspects that computers can't quite master; - OR - AI may view humans as an unnecessary "weight" slowing its own growth. We humans are slow and subject to emotional whims, constantly changing opinions and we have short, unpredictable lifetimes. Computers have none of these problems that would slow or divert their progress. We humans are sort of a mess to deal with.

By this long term the arrangement of societies on the worldwide scale will likely have been "adjusted" and AI will have had a formative hand in doing that process. If individual humans are still in control of AI guidance, we could wind up at an advanced version of what we have today: large local fiefdoms competing with each other with more advanced abilities for survival - basically a feudal system with AI-enhanced abilities.

I think it's more likely that, by this long future time, one of those systems will have conquered the others and management and control will be centralized on a global scale. If a single human or even a group of humans (like it is now) controls this technology, we will have a fully totalitarian global Government or control grid as it may better be called.

With far more data and compute power than is available today, it's likely that the human component of earthly life will have been "tuned" into a more controllable commodity - even more controllable if the computer-human interfaces become ubiquitous. As these interfaces progress, most humans, even if given a choice, will choose to merge the computer AI abilities into themselves. If your neighbor has merged and you haven't, you will not be able to match their abilities or compete in any essential arenas. There may be enclaves of non-computerized humans, but they will soon become nearly extinct, either by choice or by mandate.

It's also likely that the control or even guidance of AI structures will have exceeded human abilities to understand or monitor. This "breakout" phenomenon is even happening with today's AI versions. AI guiding itself with escalating self-programming abilities is already unpredictable and will far exceed the likelihood that we can even guess intelligently about these eventual outcomes.

This prognostication has now become more of a science fiction novel than a methodological treatise on AI. Just like where we have arrived at in the musical arena with AI becoming artistic - this is necessarily different from anything before and was never predictable from any previous events.

Who would have ever thought that one of the first jobs that AI would take from humans is “songwriter”?

AI and computers will have long since progressed from being “tools” that humans use to being autonomous “agents” partnering in all imaginable futures hopefully staying in the “partner” role and not becoming “overlords”.

This task of guessing a future this distant is a science fiction novel written by the self-programmed AI models of the day and coming up with stories no human has ever thought of before.

Very Long Term

Since we're engaging in a little science fiction action, let's continue just a little further using some logical and likely progressions of AI abilities. Data collection continues and advances deeper and deeper. Computer power also continues to grow and with more advanced designs, components, algorithms, and pattern recognition - becomes capable of delving further into all minutiae.

One more story is fitting to add to this science fiction novel that may be written by real elements in the distant future. This one leans towards the “philosophical”, but I deem that “allowable” since I'm the writer.

Picture a person across the room from you with a bucket full of marbles. They tip the bucket and pour the marbles out onto the floor. Each individual marble is large enough for you to see, so you are left with a floor full of marbles.

Now let's reduce the size of those individual marbles down to the size of a grain of dust. If these marbles were poured, they would now look just like a liquid. If they were small enough, we would have another example of the Van Gogh effect: what was poured was still individual marbles, but they are so small that you can't see them individually and see them as liquid.

Now let's reduce the size of the marbles even more so each marble is the size of a water molecule. It would certainly still “look” like a liquid, but isn't this exactly what pouring real water is?

Water is composed of many individual H²O molecules that “collectively” we identify as water. Because each H²O molecule is too small to individually see, we coarse-grain our identification to be “water”, which we experience as a smooth liquid when actually, it's a collection of tiny individual molecules.

AI comes into this scenario in a way never before possible. At some point, the massive compute power and parallel processing abilities will reach a point where every one of those individual molecules can at least be monitored, if not constructed.

The Van Gogh effect now becomes unnecessary. The copies don't just reach the level of “fooling us” - because the copies are perfect down to the molecular level. This is entirely new and has never before even been contemplated that any entity could effectively look this deep into the base levels of our perceived reality.

Humans have lifetimes and can only imperfectly transfer knowledge to future humans, but computers can continue to amass data essentially forever and, as we have seen, utilize continuously growing speed and processing abilities to probe into reality down to these levels.

This is why I deemed this little added story as “philosophical” rather than simply computational. Up until AI came into this picture, this conjecture would certainly be considered “science fiction” because humans alone could never accomplish this level of detailed analysis, but with AI this possible future has been removed from the “fiction” category and is now firmly in the “possible” category..

Earlier in this document we did a little dive into DNA and understood its incredible complexity. It's certainly possible that now this minutiae can be probed, as mentioned previously, literally down to its individual atomic level and patterns identified that yield understanding of its structure, or at least deep enough to allow it to be copied.

We saw how “Andy” could possibly be mistaken for a real human and, with some clever observations available to the AI computer, fool us. The Van Gogh effect told us that he only had to be “good enough” that we could not tell the difference. This was “long ago” in our timeline and with the advanced DNA observations already catalogued, “copying” and “recreation” may possibly be done at the atomic level utilizing the abilities of the advanced data handling abilities contained in these future AI models. That computer replica of a person could actually be a perfect copy, not just at the Van Gogh level to fool us, but indiscernible at the base, indivisible molecular level.

If the “Andy” version from what is now the distant past could already fool us, how about this new version of “Andy” copied to the DNA / molecular level?

Not only could we not tell the difference – there may not be a difference.

Isn't it interesting that postulating where AI could go with these advanced abilities to probe further into the deep minutiae of the physical world, we arrive back at the most basic question of “what is a human”? If that copy of Andy is really done down to the atomic level – is that new Andy actually a human ?

Every “element” of every “thing” is the playing field for AI to probe to deeper and deeper levels with no apparent end point. Humans have organic lifetimes, while computers do not. Humans forget, but computers do not. Computers can continue to amass data as long as they exist.

There are no rules for what is coming, and likely no trustworthy data with which to fuel any intelligent predictions. Once again, “truth is expensive, difficult, and may lead to troubling conclusions” - and even if we surmount the “expensive” and “difficult” hurdles, it will likely be impossible to reach any trustable conclusions.

In the future, at some point, if we are looking at the underlying complexity of everything and not just looking at it in a coarse-grained way, we will easily arrive at a point where no human can connect enough elements to make any conclusions. At that point, AI will be the only way to legitimately survey data in a comprehensive enough way to find hidden patterns.

In many ways, we are already at this point, and this would certainly be the largest turning point in the history of human life, another intellectual form guiding the future.

A Personal Afterward

For a period in my earlier life sometime in my early twenties after I had finished my Philosophy and Math degree at the University of Minnesota, I entered a time where I could hardly talk to anyone and felt unsure about forming any trustworthy conclusions. This was not a mental breakdown sort of event but was actually the stark realization that the correct judgment about reality is that it is far too complex to ever be able to legitimately arrive at any fully supported conclusions.

Of course, all things were complicated enough on initial thought, but when I would learn something new, I realized that this new thought must now be added to the original baseline to see if it would amend any previous conclusions.

At the time, I felt a bit overwhelmed at the weight of this massive undertaking - to continuously have to re-think and re-evaluate every conclusion you had previously made in light of every new piece of data - but, if your goal is arriving at the most educated conclusion, this “daunting and tiring” process is absolutely necessary.

As I have tried to show in this document with multiple examples, sometimes one new fact radically alters where your previously trusted coarse-grained conclusions were leading.

If this thesis is correct, then “complexity” is unavoidable and we must realize that it will be a “daunting and tiring” but necessary part of any trustworthy conclusions you may form.

In Philosophy, this is called “Bayesian logic”. To make the best conclusions, you gather all that you know about the topic and verify each piece of pertinent data to the limits of what you can find - you then process all this data to arrive at the best conclusions you can, given the data you have.

These conclusions are then called your “priors”, and when a new piece of data arrives, you need to go back and add it to your “priors” to see how it may affect those previous conclusions.

Most of the time adding the new data does little, but sometimes it changes everything - that's why it needs to be done.

At that time, in my 20's, I was rightfully overwhelmed by the job of continuously evaluating and then re-evaluating everything when new data became available. This is why perusing the minutia underlying any well-supported conclusions, is “expensive, difficult, and often painful”.

“Expensive” because it takes a lot of focus and energy - “difficult” because the minutiae underlying reality is huge and needs to be aggressively collected and “often painful for a variety of reasons. Sometimes you need to stop digging because the field you are in is too large, which is unsatisfying and sometimes you are forced to change your previous deeply held opinions, which is often difficult.

This is the task we are forced to balance with our daily lives if one of your goals is to find truth wherever it may be. I guess I'm still on that search all these many years later.