

Artificial Intelligence, National Sovereignty And War-Fighting Capability

The Rev. David R. Graham
Adwaita Hermitage
25 May 2018

This essay is a spread of observations launched against idolators, those who claim for something status it lacks, who claim ultimacy for something not ultimate, who seek to extract money and prestige from the ignorant, the gullible, and the mendacious. The essay compiles my observations and those of others, who will be indicated but not named.

The occasion for this essay is an article purportedly by Henry A. Kissinger (yes, that Henry A. Kissinger!) for *The Atlantic* of June 2018 and titled [*How The Enlightenment Ends*](#).

The decisive feature of Artificial Intelligence (AI) is its artificiality. *Artificial* means *handmade* (Latin *ars + facere*).

A truism of philosophical experience through the *millennia* is that human intelligence – call it natural intelligence (*logos*) – and primal intelligence – call it divine intelligence (*Logos*) – recognize one another because they share a common structure and power in equal measure, although the power of natural intelligence is crippled, more or less, by the dialectics of its embodiment even as that power

is cushioned, more or less, by opportunities presenting in those dialectics.

A companion truism is that natural intelligence and divine intelligence are pro-active. They are doers.

Natural intelligence is thrown into existence by divine intelligence, which thereby cripples and cushions itself but remains itself inside its crippled/cushioned existence (dialectical embodiment). As both unconditioned and conditioned, *Logos* remains plenary.

AI is a handiwork of natural intelligence, which is divine intelligence crippled and cushioned in the dialectics of embodiment. Therefore, AI is at a remove from the origin of existence in *Logos*, in divine intelligence and act. Whereas natural intelligence is thrown into existence by divine intelligence as itself plenary, AI is not thrown into existence by natural intelligence as itself plenary. AI is pieces of natural intelligence, not the whole. It cannot be the whole because its crafter is crippled.

AI is not natural intelligence and never can or will be because no thing makes itself. No existence makes its own existence. No thing has aseity. Anything handmade is less than the hand that made it. Natural intelligence is crippled yet cushioned divine intelligence but artificial intelligence, ontologically, is less than natural intelligence. The condition of its origin puts AI at permanent remove from divine intelligence and natural intelligence. It can replace neither nor substitute for either.

AI is a tool, a really capable pencil. Hooray for that, but a pencil it is and nothing more.

Leftism is a personality disorder marked by the presence of sententious spite, shameless immodesty, aggressive self-promotion, ruthless absolutism, and relentless lying. These faces leftists call democracy.

The Kissinger essay in question is tarted up leftism. It overlooks or minimizes the fundamental factor in human affairs: geography.

Who controls geography controls communications. Who controls communications has the deepest and largest effect on destiny.

The basic purpose of war and geo-politics is to control full-spectrum geography: space, air, morale, sea and land (ether, air, fire, water and earth, respectively, the five elemental principles). So-called *cyber* is fire, which is to say, electricity, heat, moral geography.

The communications nexus of the Near East (as seen from Rome) is the west coast of Greece: [Actium](#), [Lepanto](#), [Navarino](#). The communications nexus of the Far East (as seen from Rome) is the Caspian Sea, and derivatively, [Nagorno-Karabakh](#). Saudi geography, being peninsular, could not be more contingent. Still Saud is demonstrating perspicacity in recognizing that [Europeans' hatred of Jews](#) informs their support of Iran against Arabs.

First take: other than absolutely not understanding the Enlightenment and therefore absolutely misrepresenting it and *religion* (aka Christianity) . . . and other than Kissinger's age and background precluding his ability to write such an article . . . and other than his and his ghost's using this article to troll for clients of his *consulting* business . . . and other than his and his ghost's apparent unfamiliarity with [Jacques Ellul](#) on technology and history . . . not to mention Amos, Jeremiah and Isaiah . . . and other than Kissinger's being a certified, life-long anti-USA-sovereignty activist (thus the respect Democratic Stream Media and RINOs – Rockefeller Republicans – pay him) . . . and other than his being a Jew with a notorious zipper problem where *shiksas* are concerned (the same impulses Harvey *I-did-not-invent-the-casting-couch* Weinstein cultivates) . . . so far so good

As antidote, mind cleanser, I commend [How The Scots Invented The Modern World](#).

Never accept or believe *scientists'* pronouncements, especially regarding technology, and especially not if they portend an aura of menace, until you have scrubbed them against your own experience and the expectation that scientists beg funding at taxpayers' expense with less sense of shame than ever even clergy have done. Ditto for emits from the mouths of politicians and bureaucrats, especially those in office for more than four years.

Technology comes and goes. Character comes and grows.

Cardinal Wolsey was aware of his dishonor. Kissinger is not. Whatever he says is a lie. There is indeed a set of strategic and policy questions mandating decisions regarding the employment and deployment of AI. They may be orders of magnitude larger than questions and decisions regarding the employment/deployment of, say, railroads. But they are not in kind different. No work of man's hands escapes the contingency of those hands. And none, therefore, is to be feared or suppressed or presented as a threat. A creature does not supersede his creature-hood, much less can the work of his hands do. Technology is made for man. Man is not made for technology.

Augustine had many aphorisms and I have always loved this one: *credo ut intelligam*.

Traditionally translated: *I believe [in Jesus as the Christ] in order to understand [The Bible, Christian Theology]*. Perhaps better: *I accept that the structure of my being is the structure of Being Itself. Or, I participate in something in order to gain authority to think about it.*

The phenomenology of *credo ut intelligam* keeps one from a schizophrenic personality fissure.

The knucklehead young punk who wrote this article for Kissinger relies on a straw man argument (cool modern Technology bashing boring ancient Enlightenment, of which there were several, BTW, distinguished by country of origin, including a USA one) that itself depends on a false premise, namely, that some Age of Science overthrew some Age of Faith. See Augustine, above For that matter, see the great Medieval

Scholastics, Scotus, Anselm, Ockham, Aquinas, the Mutazilite Muslim Averroës, and the rest.

I have observed this scenario many times over my career as well as in the historical record: techies promising salvation and conjuring fear over something they themselves already use to harm persons, organizations, and nations. And I reprobate myself for still letting the scenario anger me. The ancient problem of idolatry, enthrallment to the work of one's hands and, most maddening to me, demand that others subscribe to one's enthrallment . . . upon pain of death actual or virtual. It raises the Elijah-like blood.

It is correct that, like any tool, AI has military and national sovereignty uses. Basically, it is a new weapon. Whether and how fundamentally transformative it is – or, in the face of intellectual inertias, could be – of diplomatics, finance, combat formation, procurement, theatre strategics, and battle tactics, etc. is an array of legitimate questions.

Already AI is weaponized by private industries, to include media and, had they their wish, elements of the IC, as well as by schools, NGOs, political parties and candidates. So it makes sense to study the statecraft asset assistance aspects (in diplomatics, finance, war-fighting) of this new weapon, simply as a weapon system/tool and nothing more intimidating or mysterious than that.

The British – [Liddell-Hart](#) – invented the concept of tank-based Lightning War (Blitzkrieg) after WW I, but the Germans – Guderian,

Rommel + – studied and applied the concept to unit composition, deployment postures, theatre strategics and tactics, and procurement.

Perhaps a cold-eyed view of AI as a capable weapons system is the wise counsel at this time regarding it. AI is not infallible, not insurmountable, not indomitable, not terrific, not omnipotent, but it is capable. It is a weapons system presenting opportunities to be explored and exploited, thoroughly,. And, into the bargain, reject anyone whose patter regarding AI presents an aura of menace, mystery, or money grubbing.

This question of AI's use begs impartial examination and assessment of problems and possibilities AI presents, and all that with a view to fulfilling the fundamental duty of national government, which is to maintain the sovereign independence of the nation and the personal independence of her citizens.

GOA MacArthur used to say that his USAAF bombers were *forward artillery* and that his line of battle moved forward behind his line of bombers. AI – and really all of *cyber* – is artillery farther forward even than the nation's bomber line. *Cyber* has been made a combatant command but its nature and utility are essentially that of fires, artillery. It is electricity, after all, which is fire.

Cyber cannot hold territory. But it can destroy morale inside a territory that is sought. Thus, like a bomber line, AI can clear the path ahead of surface (land and sea) combat forces, which can and are meant to hold territory, actual acreage.

The essential task of war is to gain and hold territory the holding of which has been deemed essential to the strategic objective and interests of the nation.

Besides outright attrition of an enemy's strength by diplomatics, finance and maneuver of surface (land and sea) combat forces, an essential means of gaining and holding territory is the destruction of an enemy's morale. That is the purpose of fires, of artillery, to include tubes, rockets, bombers, and cyber: destroy an enemy's morale, make them lose the will to fight.

Existence rises and reposes in geography.

AI is a weapon for destroying an enemy's morale from a distance and at great depth. Napoleon: *In the difficult art of war, the morale is to the material as three is to one.* In diplomatics, AI is a weapon for spreading confusion. In finance, it is a weapon for spreading distrust. In war-fighting, is a weapon for spreading chaos.

AI is a weapon of statecraft alongside others. It is useful to operators in in each of the three assets of statecraft: diplomatics, finance, war-fighting. AI has no aseity. Someone has to tell it what they want it to do. AI is a tool for helping humans achieve an objective they, the humans, define.

No tool is going to live our life and make our decisions for us. Tools help us execute decisions we have made. They are handiworks we make to make life comfortable, noble, and delightful . . . and to make more handiworks.

The young and other excited ones imagine AI can solve the world's problems. Fine. Here are some thoughts on facilitating their crash into reality whilst nursing that pretentious delusion.

BLUF:

Basically, send the young 'uns upstream to try their hand at making an algorithm to overcome the presence of disloyalty, stupidity and/or insanity at the top of a chain of authority. And have them prove by outcomes that the algorithm actually works to that effect and that they can convince top authority to use it against what they, top authority, consider their best interests, i.e., perpetuation in power of themselves, who operate by disloyalty, stupidity and/or insanity.

An interesting side-benefit to this project – to apply AI to top authority is infecting some sharp young minds with problems that will get them thinking beyond the poli-sci nonsense they suck down every day

That is true, but the reverse is true: sharp young and ambitious minds infecting top authority with ever more arcane and *necessary* poli-sci *nonsense* = lucrative contracts. Every sword has two edges.

An algorithm to prevent or override ignorance and/or cowardice (*increasingly subjective risk analysis*) among non-top authority, much less stop authority, is a wonderful mirage but a mirage nonetheless. And

having such divorced from geographical specifics so that it is applicable in any geography is a abstraction of an abstraction.

Nothing is divorced from geographical specifics and nothing appropriate for one set of such specifics is directly applicable to any other set, although fundamental grand strategic goal, with high but not absolute certainty, is constant, or should be, over all geographies at all times: namely, freedom of movement for trade and its protection there-in and there-over by nation states engaging in the movement of that trade.

On the other hand, perhaps the mirage can be used to lift the sharp young minds above the poli-sci nonsense, which also is sedition.

For example, in the field of statecraft, one could task the young smarts with creating an algorithm that illuminates paths forward for (1) rational/natural diplomatic, financial and war-fighting stewardship of a nation's sovereignty in the absence of enunciation by said nation's top authority of grand national strategic objective generally and rational grand national strategic interests by region of the globe (geography) and/or (2) the presence of open and clandestine activities by political, non-profit and bureaucratic personalities (e.g., in D-R UniParty/MSM, CIA, FBI, USAID, OECD, Rockefeller, Ford, Soros, Steyer, US Administrative/Deep State) which frustrate or degrade national sovereignty and/or national self-confidence.

That would set the cat among the pigeons. That is, the cherubic smarties would hit the wall of failure, which is redemption hard afoot. It would compel them, at least some of them, to ponder matters well above poli-

sci objectives/seditions, which are manipulation of electorates and elections. And it would be just the sort of thing for which true schools and NGOs once were notable.

Risk aversion rises with information profusion, and especially so with policy ambiguity thrown into the bargain. Let AI boys run their legs against that debility, to overcome it.

As illustration, and again in the realm of statecraft: the fundamental decision of whether the risk of a proposed action is worth the lives of subordinates – i.e. the contribution of the operation towards reaching a coherent strategic end-state is worth the cost in blood and treasure – is increasingly challenging and subjective.

Yes, challenging because the absence of rational (von Clausewitz), which is to say natural, grand national strategic policy and therefore rational, natural grand theatre strategic policy makes decisions to engage or not challenging and subjective. Meta-random might be the proper systems description. Alice at her croquette game with a flamingo for a mallet and a hedgehog for a ball.

So, where these smart ones are concerned, send them *upstream* and let their Silicone Valley cohorts accompany them informally. Give them the algorithm challenge mentioned here.

Task the young smarts with creating an algorithm that illuminates paths forward for rational/natural/geographic/organic diplomatic, financial and war-fighting stewardship of the national sovereignty in the **absence** of

enunciated grand national strategic objective generally and grand national strategic interests by regions of the globe (geography) and/or in the **presence** of open and clandestine activities by political, non-profit and bureaucratic personnel to frustrate or degrade national sovereignty and/or national self-confidence.

Underneath and over all of this is the phenomenology of freedom. Everything, really, is about freedom, which, at root, is freedom of movement, geographical freedom.

Geographical freedom is the source of spiritual freedom, which is the source of intellectual freedom, which is the source of moral freedom, which is the source of political, commercial and social freedom. Religion, too, true religion, is entirely, 100%, about freedom: finding, achieving and maintaining it as movement across geography.

Basically, send the young 'uns upstream to try their hand at making an algorithm to overcome the presence of disloyalty, stupidity and/or insanity at the top of a chain of authority. And have them prove by outcomes that the algorithm actually works to that effect and that they can convince top authority to use it against what they, top authority, consider their best interests, i.e., perpetuation in power of themselves, who operate by disloyalty, stupidity and/or insanity.

The smarties will not be able to do that because their algorithms, having no everything-ness, no eternal-ness, have no inclination of choice without inputs of current knowledge/experience. Algorithms cannot make decisions without inputs. They have no direct experience of the objects

of their inquiries. Their *experience* is entirely mediated. Humans, on the other hand, have direct experience of anything to which they turn their attention and they can make decisions without inputs of current knowledge/experience.

Awareness of that limitation of AI and contrast of it with human potential would drive the smarties, *with luck*, to develop a *lingua franca*-type re-education program, firstly for themselves, then for mid-level authority. In other words, the deficiencies of AI/algorithms must be experienced firsthand and that will show the smarties the actual problem and, *with luck*, impel them to seek actual solutions, which involve re-education on a grand scale.

... *with luck* ...

There is the situation where a tool that cannot do a job one is told to make it do compels one to discover what the job actually is and make or at least think through making a tool to do the job that is actually present. Lower authorities' re-education would follow naturally on that process were it to proceed optimally. Then upper and top authorities, though less intensively and extensively.

As responsibility increases, experience increases but expertise decreases.

First the smarties must discover on their own that their sense of mastery with AI/algorithms is mistaken. Some if not most, upon that revelation, will double-down on stupid, demand *more cow bell* . . . and money. But some will do the work (essentially theological) necessary to expose the

conundrum – that their AI/algorithm god is powerless where it matters most – and then see and spread the good word regarding how to proceed for the nation's/state's/town's/company's best interests, having been compelled to that mission by their own labors to think the whole matter through, as a Gestalt.

Idolatry is a nasty business. Tremendous power is required to smash it. But, said power is available and said smashing is indicated.

A pedagogue might suggest that focus of re-education is most productive in re lower orders of authority than upper and top.

The difference between the CIA and the FBI is that FBI catches bank robbers and the CIA robs banks.

Another Writes

- I don't necessarily disagree with his observations on the impact of the internet & general data-overload, but he omits the up-stream roles & responsibilities of proper parenting and education to develop good character. Misplacing or over-emphasizing cause and effect dynamics.

- He doesn't try to specifically dig into implications of AI to national strategies / statecraft / warfare... This is an arena that every defense contractor and techie lobbyist in DC is frothing at the mouth trying to convince DOD and politicians that their algorithm is the answer for better decision making. Gets to the discussion about DOD (especially SOF), the IC, and defense policy types having an infatuation with "Skynet" systems

& operating concepts - if only we had enough data feeds, strike forces, and this groovy algorithm we could go kill people before they even knew they needed it! Yay! (slight exaggeration...).

- A question here, which I've been thinking a lot on in slightly other contexts, is whether the domain that AI operates in could cross over to that of statecraft and warfare? Obviously in very tactical terms (weapon systems, data fusion for better logistics, tactical intel, etc.) it could, but in terms of advising or shaping national policy, grand strategy, etc? I don't see it... At the end of the day it is still making game-type adjustments based on data inputs, even if (as the article suggests) the AI could establish its own *end states*. Related question - how does one explain to the layperson that AI, no matter how advanced and able to mimic consciousness, cannot attain actual consciousness and operate in the spiritual realm? Does this inability matter practically - and why?

- He leaves out discussion of AI impacts on the financial system... or rather their rapid trading / crashing / whatever the markets and screwing us over. I'd hazard this is a bigger problem (since algorithms already trade without human inputs, to improve speed)... the system is already rigged and fragile wrt this kind of activity.

- I suspect he misstates some key points on the Enlightenment...

- In general, I think current academics (especially political *scientists*) always try to over-complicate and instill a sense of fear... and therefore a need for their continued services. I say this as a self-check; while I think a lot of worthwhile points & questions are raised in the article, my

suspicion is it isn't quite that complicated a question, or the questions are misplaced, or the motive behind the paper is deceptive. However, Kissinger's more recent work doesn't strike me as outright progressive internationalist/globalist, and he has mentioned that USA should work with Russia more (giving me some sense he isn't a political shill)... The question of his motive/agenda isn't immediately obvious to me, if it is ulterior.

This discussion on AI / machine learning is bound to come up frequently now, so it's good to have some strong philosophical background that supercharges the BS detector...

Another Writes

What immediately jumps out at me is that I don't for a second think Kissinger wrote this. The writing jumps around from reserved bonhomie (ha!) to pedantic fear mongering. I'll go a different way and comment throughout the article. A Fisking!

Three years ago, at a conference on transatlantic issues, the subject of artificial intelligence appeared on the agenda. I was on the verge of skipping that session-it lay outside my usual concerns-but the beginning of the presentation held me in my seat.

The speaker described the workings of a computer program that would soon challenge international champions in the game Go. I was amazed that a computer could master Go, which is more complex than chess. In it, each player deploys 180 or 181 pieces (depending on which color he or she chooses), placed alternately on an initially empty board; victory goes to the side that, by making better strategic decisions, immobilizes his or her opponent by more effectively controlling territory.

Any stepwise-evolving state space with a defined goal is subject to algorithmic analysis. It is not intelligence but better descriptions of the solution space that wins Go. AlphaGo is not a triumph of intelligence but of design.

The speaker insisted that this ability could not be preprogrammed. His machine, he said, learned to master Go by training itself through practice. Given Go's basic rules, the computer played innumerable games against itself, learning from its mistakes and refining its algorithms accordingly. In the process, it exceeded the skills of its human mentors. And indeed, in the months following the speech, an AI program named AlphaGo would decisively defeat the world's greatest Go players.

But it *is* preprogrammed. AlphaGo won not because of some special ability but because algorithm designers successfully defined the information the program would store in memory and manipulate on silicon.

As I listened to the speaker celebrate this technical progress, my experience as a historian and occasional practicing statesman gave me pause. What would be the impact on history of self-learning machines-machines that acquired knowledge by processes particular to themselves, and applied that knowledge to ends for which there may be no category of human understanding? Would these machines learn to communicate with one another? How would choices be made among emerging options? Was it possible that human history might go the way of the Incas, faced with a Spanish culture incomprehensible and even awe-inspiring to them? Were we at the edge of a new phase of human history?

"No category of human understanding." To me this type of statement is cover to use self-made algorithms to do whatever I want to whomever I want. And the jump to "acquired knowledge by processes particular to themselves" is similar to my saying that I've stepped up on a curb after much instruction and now will casually span the Grand Canyon.

Handwaving. And there at the end, pointing at an inflection point, "Henry" tickles us with fear.

Aware of my lack of technical competence in this field, I organized a number of informal dialogues on the subject, with the advice and cooperation of acquaintances in technology and the humanities. These discussions have caused my concerns to grow.

Heretofore, the technological advance that most altered the course of modern history was the invention of the printing press in the 15th century, which allowed the search for empirical knowledge to supplant liturgical doctrine, and the Age of Reason to gradually supersede the Age of Religion. Individual insight and scientific knowledge replaced faith as the principal criterion of human consciousness. Information was stored and systematized in expanding libraries. The Age of Reason originated the thoughts and actions that shaped the contemporary world order.

But that order is now in upheaval amid a new, even more sweeping technological revolution whose consequences we have failed to fully reckon with, and whose culmination may be a world relying on machines powered by data and algorithms and ungoverned by ethical or philosophical norms.

No, no, no. "Ungoverned by ethical or philosophical norms." No. More handwaving. Whoever writes a program, whether it is invariantly deterministic like a robotic welder or a random-annealing genetic algorithm, is completely responsible for any actions of that program. The greatest attraction of "AI" is that by stating an algorithm is so complex that "we don't understand it" and that it must be in some way intelligent we are relieved of any responsibility.

The internet age in which we already live prefigures some of the questions and issues that AI will only make more acute. The Enlightenment sought to submit traditional verities to a liberated, analytic human reason. The internet's purpose is to ratify knowledge through the

accumulation and manipulation of ever expanding data. Human cognition loses its personal character. Individuals turn into data, and data become regnant.

No, the internet's purpose is to be a multiply-redundant self-routing communication channel. Structures are built that connect to the net but they are not the net.

Users of the internet emphasize retrieving and manipulating information over contextualizing or conceptualizing its meaning. They rarely interrogate history or philosophy; as a rule, they demand information relevant to their immediate practical needs. In the process, search-engine algorithms acquire the capacity to predict the preferences of individual clients, enabling the algorithms to personalize results and make them available to other parties for political or commercial purposes. Truth becomes relative. Information threatens to overwhelm wisdom.

In other words, he and his have lost control. The clerics are unhappy, man, and they're gonna let you know.

Inundated via social media with the opinions of multitudes, users are diverted from introspection; in truth many technophiles use the internet to avoid the solitude they dread. All of these pressures weaken the fortitude required to develop and sustain convictions that can be implemented only by traveling a lonely road, which is the essence of creativity.

The impact of internet technology on politics is particularly pronounced. The ability to target micro-groups has broken up the previous consensus on priorities by permitting a focus on specialized purposes or grievances. Political leaders, overwhelmed by niche pressures, are deprived of time to think or reflect on context, contracting the space available for them to develop vision.

Finally we get to the crux of the problem. Trump! We seem to be veering from the point of AI but whatever - at least we're commiserating

on the awful fate of Our Wise Leaders, deprived by reflexive assembly of their natural reflective apprehension.

The digital world's emphasis on speed inhibits reflection; its incentive empowers the radical over the thoughtful; its values are shaped by subgroup consensus, not by introspection. For all its achievements, it runs the risk of turning on itself as its impositions overwhelm its conveniences.

"Here, let me help." We'd hear none of this had Hillary won. None of it.

As the internet and increased computing power have facilitated the accumulation and analysis of vast data, unprecedented vistas for human understanding have emerged. Perhaps most significant is the project of producing artificial intelligence—a technology capable of inventing and solving complex, seemingly abstract problems by processes that seem to replicate those of the human mind.

They seem to replicate the human mind because they **are** the human mind to the extent that a human mind copied itself to a deterministic algorithm.

This goes far beyond automation as we have known it. Automation deals with means; it achieves prescribed objectives by rationalizing or mechanizing instruments for reaching them. AI, by contrast, deals with ends; it establishes its own objectives. To the extent that its achievements are in part shaped by itself, AI is inherently unstable. AI systems, through their very operations, are in constant flux as they acquire and instantly analyze new data, then seek to improve themselves on the basis of that analysis. Through this process, artificial intelligence develops an ability previously thought to be reserved for human beings. It makes strategic judgments about the future, some based on data received as code (for example, the rules of a game), and some based on data it gathers itself (for example, by playing 1 million iterations of a game).

Reading stuff like this is infuriating. Let's say there is a complex optimization problem - multiple minima, enormous solution space, incomplete and continuing data. Much too large for a person or multiple people to solve. Impossible for a deterministic algorithm to give a trusted solution. So we make an algorithm that modifies itself as data is gathered, that searches the solution space with random perturbations to climb away from local minima and that delivers results that seem like magic. We're not too sure that they're the absolute minima but we're fairly confident they're better than anything we could otherwise find.

That is not a strategic judgement. It is not intelligence. It is the human use of a tool - computer's greater speed and memory - to attack a problem we can define but are too slow and storage-poor to solve.

The driverless car illustrates the difference between the actions of traditional human-controlled, software-powered computers and the universe AI seeks to navigate. Driving a car requires judgments in multiple situations impossible to anticipate and hence to program in advance. What would happen, to use a well-known hypothetical example, if such a car were obliged by circumstance to choose between killing a grandparent and killing a child? Whom would it choose? Why? Which factors among its options would it attempt to optimize? And could it explain its rationale? Challenged, its truthful answer would likely be, were it able to communicate: "I don't know (because I am following mathematical, not human, principles)," or "You would not understand (because I have been trained to act in a certain way but not to explain it)." Yet driverless cars are likely to be prevalent on roads within a decade.

The quandary thus proposed, to kill a grandparent or a child, would be solved by a set of rules written by a person. No intermediary can absolve the designer [from] his responsibility. Perhaps it is a thousand lines of code, or a million, or a billion; a person wrote that code and is responsible for the choice made.

We must expect AI to make mistakes faster-and of greater magnitude-than humans do.

Heretofore confined to specific fields of activity, AI research now seeks to bring about a "generally intelligent" AI capable of executing tasks in multiple fields. A growing percentage of human activity will, within a measurable time period, be driven by AI algorithms. But these algorithms, being mathematical interpretations of observed data, do not explain the underlying reality that produces them. Paradoxically, as the world becomes more transparent, it will also become increasingly mysterious. What will distinguish that new world from the one we have known? How will we live in it? How will we manage AI, improve it, or at the very least prevent it from doing harm, culminating in the most ominous concern: that AI, by mastering certain competencies more rapidly and definitively than humans, could over time diminish human competence and the human condition itself as it turns it into data.

We fear death so we Facebook, we Instagram. I'm not sure what "explain the underlying reality" means. Reality is, it isn't explained. The world is already mysterious - to really think about who we are is terrifying.

Artificial intelligence will in time bring extraordinary benefits to medical science, clean-energy provision, environmental issues, and many other areas. But precisely because AI makes judgments regarding an evolving, as-yet-undetermined future, uncertainty and ambiguity are inherent in its results. There are three areas of special concern:

First, that AI may achieve unintended results. Science fiction has imagined scenarios of AI turning on its creators. More likely is the danger that AI will misinterpret human instructions due to its inherent lack of context. A famous recent example was the AI chatbot called Tay, designed to generate friendly conversation in the language patterns of a 19-year-old girl. But the machine proved unable to define the imperatives of "friendly" and "reasonable" language installed by its instructors and instead became racist, sexist, and otherwise inflammatory in its responses. Some in the technology world claim that the experiment was ill-conceived and poorly executed, but it illustrates an underlying ambiguity: To what extent is it possible to enable AI to comprehend the context that informs its instructions? What medium could have helped Tay define for itself offensive, a word upon whose meaning humans do not universally agree? Can we, at an early stage, detect and correct an AI program that is acting outside our

framework of expectation? Or will AI, left to its own devices, inevitably develop slight deviations that could, over time, cascade into catastrophic departures?

To say that this is disingenuous is an understatement. Tay did not "become" racist any more than my walls "became" painted. Certain Twitter users quickly understood the Tay algorithm (see, I don't say "Tay's programming") and injected a new vocabulary. Simple. Nothing remarkable about it.

AI isn't "left to its own devices" to "cascade into catastrophic failures." Improperly defined algorithms with poor control of boundary conditions and lacking robust error handling will fail, as always. It's fun for fear mongers to assign human agency to the failure of algorithms sufficiently complex to appear intelligent but that doesn't make it so.

Second, that in achieving intended goals, AI may change human thought processes and human values. AlphaGo defeated the world Go champions by making strategically unprecedented moves-moves that humans had not conceived and have not yet successfully learned to overcome. Are these moves beyond the capacity of the human brain? Or could humans learn them now that they have been demonstrated by a new master?

AI does not change human thought processes. Here I agree with effect but neither cause nor example. Algorithms and processes implemented by Facebook/Google have in a decade rewired the Western mind. We don't remember many things: Google. We don't have lovely smelly odd relationships with people: Facebook.

Before AI began to play Go, the game had varied, layered purposes: A player sought not only to win, but also to learn new strategies potentially applicable to other of life's dimensions. For its part, by contrast, AI knows only one purpose: to win. It "learns" not conceptually but

mathematically, by marginal adjustments to its algorithms. So in learning to win Go by playing it differently than humans do, AI has changed both the game's nature and its impact. Does this single-minded insistence on prevailing characterize all AI?

Here, have some more fear. AlphaGo could have been written to achieve the same goals "Kissinger" says humans have but that'd be terribly boring. AI, as the term is used, will probably be designed to win. I may be going out on a limb here but if I spend a bunch of money and recruit a talented team I doubt I'll be shooting for second place.

Other AI projects work on modifying human thought by developing devices capable of generating a range of answers to human queries. Beyond factual questions ("What is the temperature outside?"), questions about the nature of reality or the meaning of life raise deeper issues. Do we want children to learn values through discourse with untethered algorithms? Should we protect privacy by restricting AI's learning about its questioners? If so, how do we accomplish these goals?

This is weak, sloppy thinking. Or, if I'm to believe that "Henry" is educated, it is condescending. Isabella learned to read in large part by a great program (AI! AI!) called Starfall Reading. We let her at it as much as she wanted - untethered, indeed - and are happy with the results. Values? Parents. Privacy? Buyer/user beware.

If AI learns exponentially faster than humans, we must expect it to accelerate, also exponentially, the trial-and-error process by which human decisions are generally made: to make mistakes faster and of greater magnitude than humans do. It may be impossible to temper those mistakes, as researchers in AI often suggest, by including in a program caveats requiring "ethical" or "reasonable" outcomes. Entire academic disciplines have arisen out of humanity's inability to agree upon how to define these terms. Should AI therefore become their arbiter?

AI as I would define it are those programs that are able to do exactly as described; to implement a well-designed algorithm more quickly and operating on a larger dataset than any person could ever hope to do. It is not, though, impossible to temper mistakes. Behind these arguments is a bit of blood, the desire to kill without remorse, to hurt people without suffering nightmares: the AI said it was okay. It makes me uneasy.

Third, that AI may reach intended goals, but be unable to explain the rationale for its conclusions. In certain fields-pattern recognition, big-data analysis, gaming-AI's capacities already may exceed those of humans. If its computational power continues to compound rapidly, AI may soon be able to optimize situations in ways that are at least marginally different, and probably significantly different, from how humans would optimize them. But at that point, will AI be able to explain, in a way that humans can understand, why its actions are optimal? Or will AI's decision making surpass the explanatory powers of human language and reason? Through all human history, civilizations have created ways to explain the world around them-in the Middle Ages, religion; in the Enlightenment, reason; in the 19th century, history; in the 20th century, ideology. The most difficult yet important question about the world into which we are headed is this: What will become of human consciousness if its own explanatory power is surpassed by AI, and societies are no longer able to interpret the world they inhabit in terms that are meaningful to them?

How is consciousness to be defined in a world of machines that reduce human experience to mathematical data, interpreted by their own memories? Who is responsible for the actions of AI? How should liability be determined for their mistakes? Can a legal system designed by humans keep pace with activities produced by an AI capable of outthinking and potentially outmaneuvering them?

How is consciousness to be defined now? These questions are off in la-la-land.

Ultimately, the term artificial intelligence may be a misnomer. To be sure, these machines can solve complex, seemingly abstract problems that had previously yielded only to human cognition. But what they do uniquely is not thinking as heretofore conceived and

experienced. Rather, it is unprecedented memorization and computation. Because of its inherent superiority in these fields, AI is likely to win any game assigned to it. But for our purposes as humans, the games are not only about winning; they are about thinking. By treating a mathematical process as if it were a thought process, and either trying to mimic that process ourselves or merely accepting the results, we are in danger of losing the capacity that has been the essence of human cognition.

The implications of this evolution are shown by a recently designed program, AlphaZero, which plays chess at a level superior to chess masters and in a style not previously seen in chess history. On its own, in just a few hours of self-play, it achieved a level of skill that took human beings 1,500 years to attain. Only the basic rules of the game were provided to AlphaZero. Neither human beings nor human-generated data were part of its process of self-learning. If AlphaZero was able to achieve this mastery so rapidly, where will AI be in five years? What will be the impact on human cognition generally? What is the role of ethics in this process, which consists in essence of the acceleration of choices?

Missing from these ramblings is the statement that Chess and Go are easily described, formally, and furthermore possess known beginning states and desired end states. Any such game or process is subject to solution by AI as described in this article. Pattern recognition, big-data analysis and gaming all share similar traits. AI is wonderful for finding solutions to defined problems in these spaces. Militarily, network analysis, SIGINT analysis, wide-area video analysis and all the other secret-squirrel stuff fall into a related problem area. AI can - and I'm sure does - work absolute magic in these and other problem areas. But it is not intelligence nor can its use absolve anyone of ultimate responsibility. To make the jump as this article's writer does many times from problems like Chess to statecraft is just silly.

Typically, these questions are left to technologists and to the intelligentsia of related scientific fields. Philosophers and others in the field of the humanities who helped shape previous concepts of world order tend to be disadvantaged, lacking knowledge of AI's mechanisms or

being overawed by its capacities. In contrast, the scientific world is impelled to explore the technical possibilities of its achievements, and the technological world is preoccupied with commercial vistas of fabulous scale. The incentive of both these worlds is to push the limits of discoveries rather than to comprehend them. And governance, insofar as it deals with the subject, is more likely to investigate AI's applications for security and intelligence than to explore the transformation of the human condition that it has begun to produce.

Sometimes I suspect that one of three things must have occurred. Either 1) all these folks with their dire warnings and rule-making are read in on information that must be classified Galactic Secret NOFORN or 2) they have a poor grasp of the underlying technology and its limitations or, rather, its definitions or 3) they think the general public/potential clients easy marks for this type of fear mongering. I suspect it's mostly (3) with a bit of (2) mixed in for fun.

The Enlightenment started with essentially philosophical insights spread by a new technology. Our period is moving in the opposite direction. It has generated a potentially dominating technology in search of a guiding philosophy. Other countries have made AI a major national project. The United States has not yet, as a nation, systematically explored its full scope, studied its implications, or begun the process of ultimate learning. This should be given a high national priority, above all, from the point of view of relating AI to humanistic traditions.

AI developers, as inexperienced in politics and philosophy as I am in technology, should ask themselves some of the questions I have raised here in order to build answers into their engineering efforts. The U.S. government should consider a presidential commission of eminent thinkers to help develop a national vision. This much is certain: If we do not start this effort soon, before long we shall discover that we started too late.

Another thought on "AI" from this morning: to be intelligent in the human sense does not mean only to perceive and order but also to do the inexplicable. Show me an "AI" that commits suicide for unexplained

reasons or that sacrifices itself in battle to save the life of a ninety-five year old blind widow, in either case unrelated to machine-learning or direct programming. Can't be done.

Another Writes

- Kissinger does not give humans credit for having an "inner necessity," which AI lacks (Kandinsky)
- Kissinger does not give humans credit for having a soul, which AI lacks, that is a part of the everything-ness/eternal-ness that is God and thereby has the inclination of choice without current knowledge/experience, which AI lacks (Boethius)
- Kissinger acknowledges that he cannot see the truth, but rather than doing anything about it, he attempts to define the problem over and over again (Plato)
- Kissinger is more comfortable with rules (ethics...the law) than freedom (morals...self discipline) because I do not doubt that he believes that human beings are inherently evil, but can be tamed by other humans (influence of, "Lord of the Flies")

I do not doubt the range of possible implications of AI, but I do not think that Kissinger made a sound argument given that he only addressed one side of the coin - strengths of AI and weakness of humans - rather than addressing strengths and weaknesses of both. Analysis of these would lead to a possible "solution" to the problems Kissinger attempted to

define, but like many, I do not think he wants to solve the problem: best to get hired into a presidential commission and think about the problem for a few years first. And then possibly think some more if the problem, 98% defined now, proves itself quite wicked.

I cannot abide the fear mongering that goes on with technological development. It is maddening. I had a little bit of fun with a visitor several months ago who gave the standard "the world is going to eat us in 20 years because of technology" blah blah blah. After reading his conclusions slide I commented that the same "wicked challenges" of the next 20 years that he outlined were the same wicked challenges that Caesar faced during his Conquest of Gaul. He was so sad. Put a big fat hole in his bubble of self-aggrandizement and perpetual pride with defining the problem over and over and over and over and over again...and over. Maddening.