China's Grand Strategy: The Importance of the Artificial Intelligence and Quantum Computing Race

Evan M. Dudik

edudik@evandudik.com

April 9, 2020

Reviewers of my recent paper, China's Grand Strategy, suggested I amend it to include an analysis of the role of two far-reaching and strategically crucial technologies, artificial intelligence (AI) and quantum computing (QC). These new technologies in the U.S.-China competition should be keeping planners up at night. My reviewers were indeed right to home in on this issue.

Let's take a closer look at AI and QC to see what they are and why they're important.

Artificial Intelligence

For the purposes of this paper, I'll define AI to include the broad spectrum of technologies which enable computers to analyze and amalgamate the judgments of human experts on large test datasets, combine their judgments, then through a Darwinian process of trial and error self-modify those judgments. From this process AI software creates a customized algorithm that examines huge datasets with the goal of equaling or surpassing the judgments of any single human expert—and examining them far more quickly than humans could do. Most AI applications report probability scores to their output.

If that's a bit abstract, here are two examples in action today. Discovery is the legal process whereby plaintiffs and defendants in a civil lawsuit must provide each other (among other things) with all 'responsive' documents. "Responsive' is what ordinary people would call 'relevant.' A lawsuit hinging on alleged price fixing conspiracy would require 'production' of all emails, presentations, and texts having to do with pricing.

However, typically, the vast majority of documents in a major lawsuit won't have anything to do with pricing. They may be marketing or human resources memos or presentations, for example. These are deemed "non-responsive." In the not too distant past, and even today, law firms hire dozens to hundreds of lawyers to pour through tens of thousands of documents to sort the sheep that are responsive from the goats that are not. This huge effort can take weeks, is subject to human error and amazing expense—even if pay for the lawyers can be as low as $25 an hour.

Today, discovery software using AI algorithms takes a small sample of documents already sorted by humans into either of responsive or non-responsive buckets, figures out what the implicit rules these experts used to make their judgments and sorts at computer speed the remaining 99% of documents. All without a bio break or sleep.

Judges trust this software—even though the algorithm developed by the software is unknown to anybody since it is developed by the software itself. One reason is that often defendants use discovery to slow down the pace of the litigation and thereby increase its expense to the plaintiff. For example, when a lawsuit involving Safeway Stores got to the discovery stage, Safeway unloaded 575,000 documents on the plaintiff. The judge mandated use of AI discovery software to cut that data dump drastically. [Casetext.com, *United States ex rel. Proctor v. Safeway, Inc.*, March 5, 2019]

A related example is how automated dictation systems first require training: the user typically reads aloud a standard text so the software can compare what it thinks the user said with the actual text. Then it modifies its natural language processing algorithms to take into account a user's intonation, timing and pronunciation. Sometimes this must be done repeatedly before the software achieves an acceptable level of accuracy.

AI CT-scan-reading software imitates the ability of expert radiologists to identify indications of lung cancer in X-ray radiographs. It's notorious because some studies indicate the software has now a better scorecard of identifying patients with cancers than most radiologists. This shouldn't be too surprising, as the software enjoys the advantage of combining the decades of experience from many radiologists.

Nothing could be more crucial than training and skill building in a technology-centric world. "DARPA [Defense Advanced Research Projects Agency], intending to reduce from years to months the time required for new Navy recruits to become experts in technical skills, now sponsors the development of a digital tutor that uses AI to model the interaction between an expert and a novice. An evaluation of the digital tutor program concluded that Navy recruits using the digital tutor to become IT systems administrators frequently outperform Navy experts with 7-10 years of experience in both written tests of knowledge and real-world problem solving." [Quoted in "Preparing for the Future of Artificial Intelligence, National Science and Technology Council," Executive Office of the President of the United States, Oct. 2016]

China's near obsession with facial recognition makes this application the country's most prominent AI focus. "[A] major recipient of AI funding in China is facial recognition. This widespread technology covers streets and most public spaces such as banks and train stations. The government uses it for everything from identifying jaywalkers to allocating toilet paper. More significantly, it's also been embraced by the government as a tool for surveillance and tracking. This is a technological advantage that US citizens probably wouldn't want to replicate." ["China overtakes US in AI startup funding with a focus on facial recognition and chips," James Vincent, *The Verge*, Feb 22, 2018].

AI isn't just a software race. Artificial intelligence can benefit from integrated circuits optimized for AI. Reason: to learn and distinguish among, for example, all the different configurations of human faces requires crunching huge amounts of data. Says Gregory Allen, former Adjunct Senior Fellow at the Center for a New American Security (CNAS):

> Many traditionally software-focused U.S. technology companies, such as Google and Amazon, have created and acquired semiconductor design divisions specifically to work on AI accelerator chips. These chips can offer dramatically superior performance over GPUs for AI applications even when manufactured using older processes and equipment. The first generation of Google's primary AI chip, called a Tensor Processing Unit (TPU), for example, is manufactured using 28 nanometer process technology, which is already widely available in China. Google claimed in 2017 that its first generation TPU was 15–30 times faster and 30–80 times more power efficient for AI workloads than contemporary GPUs.
>
> Chinese firms Baidu (in partnership with Intel), Alibaba (via a new subsidiary, Pingtouge), and Huawei (via its HiSilicon subsidiary) have all established semiconductor design divisions focused on developing AI accelerator chips. Chinese AI chip startups Horizon Robotics and Cambricon have raised hundreds of millions of $USD in venture capital funding at multibillion-dollar valuations. [Gregory Allen, "Understanding China's AI Strategy: Clues to Chinese Strategic Thinking on Artificial Intelligence and National Security," Center for a New American Security, Feb. 2018. Citations omitted].

China is investing unstintingly in AI:

> "Last summer, China's State Council issued an ambitious policy blueprint calling for the nation to become "the world's primary AI innovation center" by 2030, by which time, it forecast, the country's AI industry could be worth $150 billion. "China is investing heavily in all aspects of information technology," from quantum computing to chip design, says Raj Reddy, a Turing Award–winning AI pioneer at Stanford University in Palo Alto, California, and Carnegie Mellon University in Pittsburgh, Pennsylvania. "AI stands on top of all these things."
>
> "In recent months, the central government and Chinese industry have been launching AI initiatives one after another. In one of the latest moves, China will build a $2.1 billion AI technology park in Beijing's western suburbs, the state news service Xinhua reported last month. Whether that windfall will pay off for the AI industry may not be clear for years. But the brute numbers are tilting in China's favor: The U.S. government's total spending on unclassified AI programs in 2016 was about $1.2 billion, according to In-Q-Tel, a research arm of the U.S. intelligence community. Reddy worries that the United States is losing ground. "We used to be the big kahuna in research funding and advances." [China's massive investment in artificial intelligence has an insidious downside", Christina Larson, Science, Feb. 8, 2018].

China sees the military applications of AI as the core of an inevitable arms race. It has responded with massive funding. Says Lt. Gen. Vera Linn "Dash" Jamieson, deputy and chief of staff for intelligence, surveillance and reconnaissance on the Air Staff at the Pentagon. "We estimate the total spending on artificial intelligence systems in China in 2017 was $12 billion. We also estimate that it will grow to at least $70 billion by 2020." ["China Leaving U.S. behind in Artificial Intelligence," [Mariana Pawlyk, Military.com, July 30, 2018]

AI's military applications include surveillance of enemies, putative and real terrorists (e.g. disfavored minority populations), autonomous vehicles and autonomous lethal weapons. In fact, China's government believes the country's culture and strategic position give it advantages over the US. Certainly, implementing AI is easier to implement because privacy is the least of the CCP's concerns. This allows China to develop obscenely huge facial and demographic databases and learn more quickly than the U.S. how to advance AI technology, both hardware and software. AI "…will be easier to implement in China than in the United States," says expert Gregory Allen [Gregory Allen, "Understanding China's AI Strategy: Clues to Chinese Strategic Thinking on Artificial Intelligence and National Security," Center for a New American Security, Feb. 2018].

Is AI "just another" leading edge technology? Why should this paper call out AI among so many others, such as hypersonic missiles, underwater drones, anti-satellite satellites? The reason for calling out AI (and for that matter, quantum computing) is that AI embodies the capability to leap-frog U.S' current military advantages. AI is an enabling technology that super-charges other emerging military technologies such as swarming hypervelocity space, air and underwater missiles and torpedoes. China believes the U.S. is so wedded to its $10 billion a copy aircraft-carriers and their aircraft-centric weapons platforms that it can't adapt to a new world of AI-targeting, AI-guided, AI-allocated, swarming autonomous weapons systems.

Allen again:

> "…China [believes it] is better positioned to adopt military AI than the United States. In this theory, the United States' current advantages in stealth aircraft, aircraft carriers, and precision munitions actually would be long-term disadvantages because the entrenched business and political interests that support military dominance today will hamper United States in transitioning to an AI-enabled military technology paradigm in the future. As one Chinese think tank scholar explained to me, China believes that the United States is likely to spend too much to maintain and upgrade mature systems and underinvest in disruptive new systems that make America's existing sources of advantage vulnerable and obsolete…."
>
> "Zeng Yi, a senior executive at China's third largest defense company [says] '…In future battlegrounds, there will be no people fighting.' Zeng predicted that by 2025 lethal autonomous weapons would be commonplace and said his company believes ever-increasing use of military AI is 'inevitable.'" [Allen, op. cit.]

-

AI, however, isn't foolproof or spoof-proof. For example, one weakness of facial recognition AI is that the software generates its own, highly complex algorithm. It is so complex, as you might expect from what's essentially a trial-and-error evolutionary approach, that it's generally impossible to identify exactly what the algorithm is. (This is true of most or all other AI applications, too). Further, any pollution (intentional or not) or incompleteness in the teaching set provides an opening for incorrect assessments. Sometimes what AI image-recognition software identifies with high certainty are drastically mistaken: Sea turtles for rifles, panda bears for monkeys and distressingly, for believers in autonomous vehicles, stop signs with a few bits of masking tape for 45 MPH speed limit signs or a green light. [See, among many articles, "Slight

Street Sign Modifications can Completely Fool Machine Learning Algorithms." By Evan Ackerman IEE Spectrum Aug 4, 2017; "Attacking Artificial Intelligence: AI's Security Vulnerability…", by Marcus Comiter, Belfer Center, Harvard Kennedy School, August 2019].

China is developing AI to spoof AI. Software that can learn to identify military targets in satellite photographs can also learn how to disguise those objects. Just as in face recognition, AI depends on identifying unique shapes and their angles and distance relationships to one another. These shapes can be altered or disguised. It's possible, though awkward, to do the same to your own face.

And, you guessed it: Malicious AI can be used to learn how security codes inside AI algorithms themselves work, and work around these security measures. This is above all a concern to financial institutions. The worry is that AI mole algorithms could infect global banks' computer system as moles and take over when the time is ripe. [Michelle Cantos, "Breaking the Bank: Weakness in Financial AI Applications," Fireye.com, March 13, 2019] No imagination required to understand how a malicious state actor, perhaps ones beginning with the initials C or R, might be tempted to invest in such a disruptive capability.

Quantum Computing

For the purposes of this paper, I'll define quantum computing as the use of the hard-to-fathom properties of subatomic particles to provide more powerful logic for solving exponentially more difficult problems than can be solved by today's 0-1, 2-bit computing logic.

QC machines trap and isolate subatomic particles in a state of limbo with respect to one of several properties, such as 'spin'. These properties become resolved only when the particle is observed (although I have yet to come across a decent definition of 'observed'). The in-between or both-and state is called 'superposition.'

QC also uses the quantum of property of entanglement. Entanglement is the behavior subatomic particles manifest of being in the specific condition (exactly the same or the exact opposite) as a distant twin particle. When subatomic particle A switches state (e.g. its 'spin') its twin, particle B switches to the opposite state (e.g. spinning in the opposite 'direction', instantaneously, no matter how far away it is in space—and possibly, time.

Entanglement's importance lies in the fact that compared with conventional 0,1 bits quantum bits, or 'qubits' in superposition, if entangled, provide more than 2 bits of information; they supply 4 bits, vastly multiplying computer power in special situations.

Quantum computing doesn't just speed up what a classic computer can do; it can solve problems that in principle a class computer can't solve. For instance, some computational problems involve more potential solutions to the problem than there are subatomic particles in the universe. An oft-used example is charging through an unsorted database, say a phonebook, searching for the name that corresponds to a given phone number. If the phonebook is very large, this will take a classical computer a very long time. A quantum computer can find the match lickety-split. Another example is identifying the longest (or shortest) route among a traveler's

multi-destination trip in which there are a very large number of possible combinations of routes. (Note: there seems to be a consensus, not unanimous, that QC can solve this famous "Traveling Salesman" problem. Here I'll go with what appears to be the majority opinion).

The application that's frightening to security experts—is the ability of quantum computing to identify the prime number factors of any integer no matter how large. Today's best encryption tools rely on the practical impossibility of classic computers to accomplish this even given a lifetime.

Quantum computing has a close relative in quantum communications. Here highly secure encryption is the quantum device's goal. In August 2016 China inaugurated the first space-based quantum-secured communications link. A satellite called Micius (named after an ancient Chinese astronomer) provides the power to send entangled photos to two different Earth stations. The entangled photons create quantum encryption keys to the Earth stations. This creates secure communications since any interception of the codes causes the super-positioned photons thus observed to collapse, triggering notice that the communications have been compromised. In September 2017 the Chinese Academy of Sciences hosted a quantum-secured videoconference with researchers in Vienna, Austria. ["Is China the Leader in Quantum Communication,"www.insidescience.org;Yuen You, Jan 19, 2018; "Why Quantum Satellites will make it harder for states to Snoop", Jacob Aron, New Scientist.com, August 24, 2016; "China's quantum satellite sends unbreakable signals from space," Financial Express, August 10, 2017]. This might make the U.S. billions of dollars' worth of eavesdropping equipment used by the National Security Agency and others obsolete.

China's quantum communications efforts haven't stopped at space-based communications. It has demonstrated that a fleet of quantum-entangled photon emitting drones can make a localized network interception free. ["China is developing drones that use quantum physics to send unhackable messages, "Stephen Chen, South China Morning Post, Jan 10, 2020]

Unlike some of today's AI applications, QC isn't ready for prime time. There remain two big obstacles. First, the physical—that is, engineering, roadblocks. The most frequently used approach to creating and holding quantum particles in superposition, the approach IBM uses, for example, employs extreme isolation of the subatomic particles whose 'uncollapsed' state at a temperature colder than interstellar space. Holding them long enough to manipulate the probabilities of how they will be observed is at the heart of its ability to address many potential solutions at once (although this 'at once' description is distortion promoted by the popular press). It's a true challenge to maintain the isolation from all interfering particles/waves flying around Boston or Beijing or any of Russia's 95 dedicated R&D cities. Secondly, the mathematics community is only now developing the algorithms which can put QC to work solving real-world problems. A big part of this latter problem is that QC is both noisy and inherently probabilistic. This means that typically a computation needs to be run many times before a probable answer becomes prominent while other answers become extremely improbable. As in conventional computers, QC requires great attention to error correction features—which themselves require QC power.

But the day for QC is dawning. IBM provides free access to a demonstration quantum computer which has already solved thousands of problems submitted to it online from around the world. More than 200,000 users, including ExxonMobil, Daimler, Los Alamos National Laboratory, Stanford University and Mitsubishi Chemical have used the computer. [Forbes, Mark Hunter (from IBM), "Quantum Computing is Here" Jan 20, 2020]

CFO Technology.com correctly qualifies the usefulness of at least today's vision of QC:

> …[Q]uantum computers don't deliver one clear answer. Instead, users get a narrowed range of possible answers. In fact, they may find themselves conducting multiple runs of calculations to narrow the range even more, a process that can significantly lessen the speed gains of doing multiple calculations at once.
>
> … quantum computers will be used for different kinds of problems — incredibly complex ones for which eliminating an enormous range of possibilities will save an enormous amount of time.
>
> Quantum computers have four fundamental capabilities that differentiate them from today's classical computers: (1) quantum simulation, in which quantum computers model complex molecules; (2) optimization (that is, solving multivariable problems with unprecedented speed); (3) quantum artificial intelligence, with better algorithms that could transform machine learning across industries as diverse as pharma and automotive; and (4) prime factorization, which could revolutionize encryption."

Besides IBM's super-cooled system, other hardware approaches vie for the crown. Among them are; Using silicon chips to leverage integrated circuit manufacturing technology that's already making circuits at near-quantum scale; photon-based qubits, trapped ion and quantum annealing approaches. The state of quantum computing technology appears analogous to the that of automobile technology circa 1900. At the turn of the 20th century, nobody knew whether the future fuel for cars was going to gasoline, electricity or steam. All had their adherents. A Stanley Steamer won the world speed record at 128 miles per hour in 1906, for example.

Just as the gasoline powered automobile needed an infrastructure of oil refining, fuel distribution, dealership and repair shops to take off, QC has the formidable task of developing the physical structures, calculation-and-control algorithms and peopleware infrastructure to make QC come to life. An example problem is that to provide error correction and control, quantum computers now envisioned require a vastly greater number of performing these 'overhead' functions than qubits doing actual work. (Note: The QC community, not me, uses the term, 'overhead.').

So What

I'll focus now on the geopolitical implications of AI and QC. The table below shows the broad social, political, scientific and commercial fields of human endeavor and conflict likely to metamorphize under the impact of these technologies.

Where AI and QC Change the World—Examples
Note: In many cases, AI and QC can be expected to complement and empower each other.

Domain	Sector	AI Relevant?	QC Relevant?	Comment
Governance				
	Population surveillance and terror/rebellion/dissent threat detection	Yes		A reality today
	Resource allocation for health services, education, capital investment	Yes	Likely	China already starting this
	Population movement forecasting	Yes	Yes	
	Emotion detection for compliance monitoring, lie detection	Yes		
Geopolitics and Military				
	Nation-state decision-making	Yes	Yes	Scenario building and probability assessment
	Military targeting: acquisition, identification, positioning and striking	Yes	Likely	Could assist in overcoming advantages of stealthy aircraft and naval platforms
	Reconnaissance and surveillance; intelligence interpretation and automated threat warning	Yes	Possible	Being explored now

	Adversary behavioral analysis; anomalous behavior detection	Yes		A reality today
	Military tactical and operational decision-making	Yes	Likely in specialized cases, at minimum	Scenario-building for the networked battlespace
	Weapons system design, modeling and testing	Yes	Yes	
	Training enhancement; personnel deployment	Yes		
	Disrupt global finance	Yes	Probable	
Science and Technology Research				
	Pharmaceutical research and discovery	Highly likely	Yes	An early candidate for QC
	Fundamental scientific research		Yes	
	Materials research	Highly likely	Yes	
	Technology adoption decision-making	Possible	Yes	
Economics and Finance				
	Economic/trade flow modeling	Yes	Probable	
	Financial forecasting and scenario building	Yes	Probable	Sentiment analysis; adversary/competitor behavior modeling
	Merger and Acquisition modeling	Yes	Probable	
	Investment and money management	Yes	Yes	
	Loan analysis*	Yes	Possible	
	Automated trading*	Yes	Possible	Highly vulnerable to hacking, especially by state actors
	Compliance and *fraud detection	Yes		Detect behavior anomalies
Societal Relationships				

	Legal processes	Yes		Happening today
	Education customization and outcomes	Yes		Customized curricula, testing and outcome prediction
	Land use planning	Possible	Probable	Example: traffic forecasting and network scenario-building
	Income redistribution and taxation scenarios		Probable	
Industry and commerce				
	Mineral/hydrocarbon discovery and exploitation	Yes	Yes	
	Product design	Yes	Probable	
	Product manufacturing			
	Applying technology	Yes	Probable	
	Logistics efficiency improvement	Possible	Yes	
	Customizing customer and user experience	Yes		Primitive voice recognition and contextual response deployed now. "Sorry, I didn't get that"
	Marketing and sales effectiveness and efficiency; targeting	Yes	Possible	
	Manufacturing optimization	Probable	Yes	
	Employee behavior control and analysis; training	Yes		AI for recruiting employees happening now
	Customer, employee emotional analysis	Yes		Amazon and Microsoft have primitive, racially and gender-inept versions for sale today

*["Breaking the Bank: Weakness in Financial AI Applications, Michelle Cantos March 13, 2019. Fireeye.com]

You can easily see why China's has committed to investing $10 billion for building the world's largest QC research center [Forbes, April 29, 2018]. These technologies have the notice of the Politburo of the Chinese Communist Party and President Xi Jinping who is calling for China to become: ".science and technology superpower, quality superpower, aerospace superpower, cyber superpower" and "intensifying cooperation in frontier areas such as digital economy, AI, nanotechnology, and quantum computing, and advancing the development of big data, cloud computing, and smart cities so as to turn them into a digital silk road of the 21st century." [Center for Strategic and International Studies, "Cyber Policy and the 19th Party Congress," Oct 26, 2017]. China clearly believes that these technologies could provide the eagerly sought-for technological frog-leap over the U.S. and the West.

AI, QC and China's Grand Strategy

In China's Grand Strategy, I argued that China's grand strategy follows an Isolate, Neutralize and Dominate architecture:

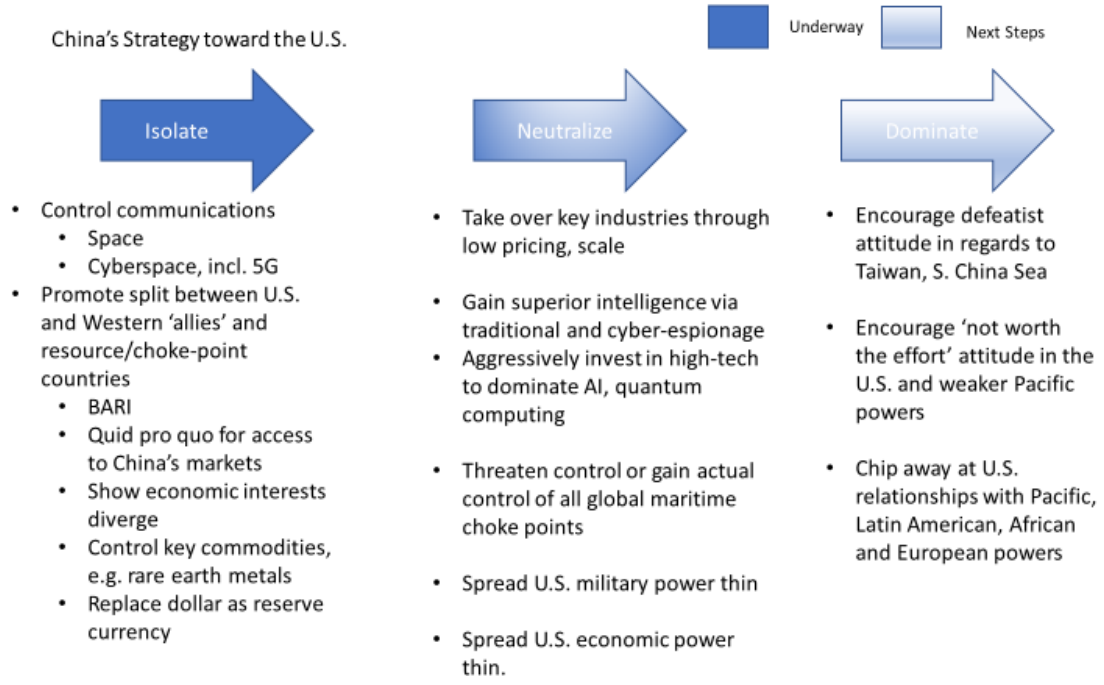

Chinese competitive investments in AI and QC support well many of its strategy's key building blocks:

Isolate:
-Controlling space and cyberspace communications: QC useful for breaking encryption codes of satellite and computer network communications

-Controlling key commodities: Identifying key minerals, energy sources and uses and controlling their transportation could benefit from both AI and QC. This could include pharmaceutical research and production
-Replacing the dollar as the world's reserve currency: Both AI and QC and help model, predict and manipulate foreign currency flows

Neutralize:
--Gain superior intelligence: AI and QC promise faster analysis of large amounts of intelligence/espionage/surveillance information—including the ability to identify and target here-to-fore stealthy aircraft and submarines
--Use QC to aggressively cut manufacturing costs, supporting China's price-cutting strategy
--Model U.S. military capabilities world wide under various scenarios using QC, enabling China to deploy its forces in a way that stretches the U.S. thin

Dominate:
-Use AI to model U.S. reactions and China's degrees of strategic freedom under various 'dominate' strategies

What is the Upshot?

There's no doubting that the U.S. is in a war with China with respect to AI. AI has too many geopolitical/military implications for this not to be the case. Quantum computing can hardly be different; just a little earlier on the adoption and implementation curve. (I've already discussed China's recent implementation of quantum-based secure long-distance and local communications).

So, if we ask ourselves what to do, we need to look at what asymmetries in our respective situations reveal.

Unfortunately, the most obvious features of the competition reveal asymmetries that favor China:

In the normal course of events, the U.S. and (some of) its allies would engage in an intellectual and technological arms race of attrition: who can make the most progress the fastest in weaponizing these technologies? Who can write superior code for AI algorithms faster? Who can miniaturize and make robust quantum qubit environments? Can the U.S. invest more money—in terms of purchasing power parity--and invest more productively than China?

In such an arms race, it's hard to see that either of the two societies has any inherent advantage. Some claim that the U.S. more open society provides a long-term advantage. I'd like to believe that myself, but the facts show that whatever may be said about Chinese society, its technological prowess in these fields hasn't been much stymied by its closed nature.

An asymmetry that cuts against the United States is that U.S. corporations which have subsidiaries in China are, by Chinese law, required to support the Chinese government's organs of state security. I noted this in the original China's Grand Strategy article. This means that Alphabet, Inc.'s Google subsidiary, which has R&D, hardware and software operations in

Beijing, must provide any and all requested information and resources to the Chinese government. You'd think that after being censored out of existence in China as a search engine, Google would have learned its lesson—but the lure of the China's huge market and the opportunity to glean some of the leavings from the table of China's AI research seem too hard to resist.

The porosity and dependence on Chinese university students constitute another symmetry that cuts against the United States in the AI/QC competition. U.S. colleges and universities host about 400,000 Chinese students per year [Statista.com, as of November 2019]. "Students from China earned 5,157 doctorates in science and engineering fields at American universities in 2017, accounting for more than 12 percent of the 41,438 doctorates awarded in science and engineering fields in the U.S." [*Inside Higher Education*, June 4, 2019]. Many study AI, QC and other STEM topics. They take the U.S.' most advanced research back to China. The Chinese government actively reminds students breathing the rather different air of Seattle, Boulder, Colorado, Georgia Tech and MIT that their first loyalty is to the Middle Kingdom.

U.S. academic institutions of course receive tuition payments from these students, often at the list prices U.S> parents can't afford, effectively subsidizing them. Chinese students fill classes that otherwise might not be taught. They populate research labs that might otherwise find staffing difficult. The downside: According to the U.S. Department of State, "The U.S. intelligence and law enforcement communities have identified an increasing number of instances in which foreign intelligence services co-opt academics, researchers and others to conduct activities on behalf of foreign governments during the individual's stay in the United States" [Inside Higher Education, June 4, 2019]. The U.S. is reviewing visa applications more carefully, issuing more 1-year visas vs. 5- and 10-year visas.

As a top-down society, China can target STEM sectors within AI and QC it wants to focus on. This has its advantages and disadvantages. The advantages revolve around the ability to quickly mobilize talent and resources and the clarity that comes from well-defined tasks. That very focus may provide a disadvantage in that other important sectors might be ignored or starved. I expect, however, that China will sufficiently resource any R&D sector for AI and QC that promises the faintest glimmer of strategic advantage. In the short-term, advantage China.

The view from 100,000 feet: China can take advantage of broader asymmetries than those affecting the AI and QC races. For example, China can use the U.S.' political and personal indiscipline and lack of self-control to push the country further and further into debt. (The fastest growing line item in the U.S. Federal budget is debt service now growing even faster with the $2.2 trillion coronavirus bailout).

This kind of broad asymmetry suggests the importance of a lever against China I alluded to in China's Grand Strategy. And that is to exploit the two-edged sword of AI to weaken the apparent monolith that is the Chinese Communist Party. In his great work, *Grand Strategy of the Byzantine Empire*, Edward Luttwak points out for most of its 1,100 year existence the Byzantine Empire was outnumbered in population and surrounded geographically by contemporary and potential enemies. Byzantium evolved several stratagems to counter-balance its negative asymmetries. Chief among these was to suborn opponents' generals and political leaders with a

range of tools ranging from flattery to bribery to even grants of large estates. This bred distrust among the leadership of Byzantium's enemies, providing a wedge for Byzantium to divide and defeat, if not conquer. China increasingly depends on AI to assist its elite to control its disparate and restless population. It's also hard to fathom that the CCP elite won't use AI to control the People's Liberation Army and one another. So, a lever for an adversary is to subvert China's AI infrastructure to sow discord and distrust. There should be fertile ground here as many of China's large cities increasingly act like local satrapies, bowing and scraping to Beijing only as necessary.

There's plenty of distrust to go around. Among the population at large, protests over local grievances are common. Local governments in cahoots with developers expropriate the land of 4 million Chinese every year, on average reselling the land to developers at a 4,045% mark-up.[Elizabeth C. Economy, "A Land Grab Epidemic: China's Wonderful World of Wukans," Council on Foreign Relations, " February 7, 2012]. I've known for many years that U.S. county-level politics is in the hands of the real estate developers, but it's not quite as bad as in China. The CCP approves of protests, up to a point. It is cagey enough to know they provide a great way for the population to let off steam. While waving their signs, protesters generally affirm their Party fealty in the most extreme terms.

It wouldn't take much to sow discord.: Just identify a few key behaviors the Chinese AI systems regard as tell-tales for disloyalty and using those to subvert their algorithms. This in turn requires robust Western intelligence and some access to either the test sets of Chinese AI social behavior surveillance or to those key tell-tales directly. This is a job for the CIA or NSA directly or indirectly through the existing dissidents, power-seekers and oppressed minorities in China.

Since quantum computing and communications remain nascent, sources of symmetry and vulnerability likewise remain emergent. We know, however, where to concentrate at least the U.S. defensive work: communications, encryption and target detection. We can only hope that the defense establishment has the recognition and the country has the intellectual horsepower to ward off these threats as it warded off the USSR's nuclear and third-World adventurism 1960-1989.