All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Some results in cybersecurity and why they may be interesting

Every once in a while I come across something interesting with substantial potential impacts but that differs from the common misconceptions. Many of them I point out with a fevered disdain of foolishness, while others I view more philosophically. This article is about some recent results that I think are worth attention, that are counter to much of what you might have heard or understood, and differ substantially from the hyperbole we see in the marketplace.

AI, Neural Nets, and Complexity of the Brain

For a long time I have known that current neural nets are nothing like decent approximations of the human brain cells. Now science is starting to speak to these issues. In a recent article about computational complexity of a single neuron[1] it was identified that the current closest computational approximation of an actual neuron by a computational model is (for 99% accuracy) “In most of the networks, that equated to about 1,000 artificial neurons for just one biological neuron.” These models involve scores of layers of artificial neural network elements, far higher than the usual 3-8 layer versions we see today.

I started studying this area in the 1970s, where I proposed a model for implementation in silicon. That particular proposal never went anywhere, but it has long been known that the operation of real neurons depends on complex chemical processes in feedback systems involving the brain and body. As such, no purely computational local (single neuron) model is every likely to do what the brain does, because the brain varies functionality non-linearly with the analog nature of different chemicals working their way through the brain, and in addition, the neurons themselves are analog devices grown in an analog environment with shapes and other features that are not digital in nature. Because the physics of digital information differs from the physics of analog (real-world) systems[2], this may also limit the ability to model with digital approaches at a more fundamental level. In addition, neural firing in the brain range over a frequency range and are analog across that range, again introducing a challenge regarding continuous (brain) vs. discontinuous (computational neural networks) mechanisms. The question of precision and accuracy of models and their effect on outcomes, error accumulation issues, discontinuities surrounding firing events, and other related matters produce many distinctions, some which may or may not make for differences between artificial and natural intelligence.

Mostly, we don’t actually know how these things work, and today, science cannot even accurately model a single biological cell of even a very simple sort. At the physics level the volume of space is too large for accurate modeling, while at the chemistry level, some parts are modeled computationally at high cost but only limited aspects of folding, and these often lead to large numbers of possibilities only a few of which are actually realized in nature. They reduce the search space in things like drug discovery, but do not provide modeling of the sort usable for emulating cellular function or neural activity.

1 “How Computationally Complex Is a Single Neuron?” https://www.quantamagazine.org/how-computationally-complex-is-a-single-neuron-20210902/

2 “The Physics of Digital Information” http://all.net/books/2013-DFE-Examination.pdf Chapter 3 (pp 83-139)

Quantum cryptography and its real effect on current systems

I hear the constant drum beat surrounding quantum cryptography ending traditional cryptographic systems and causing some sort of crisis in the coming years. But the current reality appears to be very far from that. A recent NSA document on these issues[3] may be enlightening, but all such official releases from intelligence agencies regarding cryptographic systems must be taken with a grain of salt. Having said that, I should also point out that this article reflects the current requirements for protection of government systems, so the equities issues[4] aside, their information is reasonably accurate from a standpoint of my own expertise and other 3rd party information. Here's (in extracted reformed summary) what they say:

- Today, and likely for the next 20 years, quantum computing attacks on cryptographic systems will not have a sufficiently meaningful effect to warrant changes in the current approaches other than possibly increasing key sizes, which has to be done anyway for non-quantum computing scaling of computational resources using traditional means.
- Quantum key distribution is operable today at the quantum level (basic physics) but building actual systems that do this securely is not yet demonstrated to be effective for high surety, and existing systems can be defeated even if the quantum aspects work.
- Quantum random number generation is viable today and may be used if properly implemented for the situation at hand.

For those who may be confused by these three different aspects of quantum physics and its implications to cryptographic systems, if you don't understand these differences, you should probably not evaluate, invest in, operate, or opine on these issues until you do. You will be highly susceptible to frauds.

Will training help to counter influence operations?

My personal experience tells me yes, but don't expect miracles. There has been too little scientific research on the retention periods for cyber-security training writ large, and almost none I could find on retention for countering deception. My current recommendation is at least once every 6 months for medium or lower consequences and more frequently for higher consequences (3 months or monthly with verified learning and multi-modal for high and extreme consequences). We see training requirements for things like operating a commercial airliner, but no real studies I am aware of have been done for cyber-security training, much less countering deception. The closely related counter-intelligence training regimens I have seen have 6-month requirement, but again with no basis I have been able to identify.

My 6-month period stems from the early days of computer viruses where a new virus would be effective at tricking users into actions, but once major publicity was released, the variants that followed were far less effective. Every 6 months or so, a new one would be successful and trigger another training with a reduced success rate for the attackers. Frankly, this is a poor way to make such decisions, but until we get a better way, there it is.

3 Quantum Computing and Post-Quantum Cryptography https://media.defense.gov/2021/Aug/04/2002821837/-1/-1/1/Quantum_FAQs_20210804.PDF

4 The equities issue is about the tradeoff between offense and defense. The NSA in particular is referenced in various places as wanting cryptographic systems that only they can break. Thus their public statements, and even those applied to defense systems of the US, may be understood to imply that what they say is what they want everyone to do, but not necessarily the reality they see within their classified (offensive) environments.

Blockchain, distributed ledger, and crypto-currency

These are of course three different things. Blockchain is used to implement a (relatively) high integrity distributed ledger mechanism to implement crypto-currency, but such systems also tend to be brittle in various ways. Also note that:

- Private such systems are completely different than public ones, even though they may use the same algorithms.
- Many such systems use large amounts of computing (and electrical) power because they are based on high complexity computations.
- Because the computational cost may exceed the value, many people steal the computational power from 3rd parties, in some cases unwittingly.
- Many people claim and believe these systems are untraceable, and thus some of them are widely used for criminal activities and otherwise barred financial transactions, but of course they are traceable, even if they are not traced all the time.

To be clear, I support the use of blockchain and distributed ledger and similar mechanisms in private detection of inauthentic transactions and altered storage. I think they are a good idea, and why would I not, given that I spent years trying to convince people of the benefits of integrity protection via cryptographic checksums, which is what these systems do.

However, crypto-currency is problematic for many reasons and has been demonstrated to not have the miraculously claimed benefits many associate with them. They are brittle (a change the gets into the system can cause long-tailed consequences), not good for high volume transaction systems (because they slow very significantly with volume), susceptible to denial of services (little resource is required to cause them to slow to a crawl), susceptible to takeover (pay to build N+1 members of the group when N exist and you can screw up the whole system), and of course the implementations have flaws that can cause common mode failures and subvert the systems.

Don’t trust zero trust

I have written several articles on this subject this year and you can read them online. I don’t trust zero trust because it’s (1) a misnomer (not actually about zero trust but about risk aggregation in systems and people that then have to be trusted more) and (2) most write-ups on it are pathetic in their failure to address the set of issues comprehensively. A shout out goes to Ron Ross of NIST who has done perhaps the best job of keeping his internal requirement of talking about zero trust architecture while trying to do good things along the way with trying to find ways to build more trustworthy systems.[5] But the problem is huge and the attention to it is nominal at best. Lots of money is spent on ridiculous things, but very little on the basic science and engineering discipline of security.

Some of the fundamentals that most people who follow these paths seem to ignore include complexity, granularity, and the translation between security policy and technical security policy. I will go into more depth on these below, but at a basic level, these problems are fundamental to the limits of computers and people and the way they connect to each other. As such, these challenges are unlikely to be changed by technology for the foreseeable future. As such, they are basics not apparently being taught to so-called security experts.

5 NIST SP800-160 https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v1.pdf

Basics of complexity and granularity and tradeoffs of space and time

This is actually a very old and well known issue, and one which I have written about a lot. It has to do with the fundamental problem with complexity and the tradeoffs associated with granularity. I will start with the basics of time space tradeoffs in digital circuits, an early result regarding circuit design, and one that I learned in the 1970s as an undergraduate student.

Any combinational logic equation can be implemented by a 2-level sum of products circuit. In essence, for each combination of N input variables (bits), a circuit doing an AND of all the bits that are '1' and the inverse (NOT) of the bits that are '0' can be used to generate each output bit desired. By doing an "OR" on the output of all of the possible combinations of input bits, each output bit can be generated. This can be done for each output bit. The time to produce this result is only the time required to pass through 2 levels of gates: 1 AND level (with input inversion for some inputs) and 1 OR level. Using more power, these logic gates can run faster, and 10^10 bits per second or more can run through circuits of this sort. So any combinational logic output can be computed from inputs in as little as 2/10 of a nanosecond.

But in exchange for that incredible speed of calculation, the number of circuits (space) is potentially enormous, in particular, for N input bits to produce each output bit takes up to 2^N gates, and of course for M output bits we need M*2^N gates (+M gates for the OR circuits). On the other hand, we can get tremendous space reduction in most cases by creating a circuit that is slower. For example by using a sequential adder, we can generate an N-bit addition result using 5 gates in time N. This is an example of a time space tradeoff that is fairly universal at the circuit level. I should also note that the speed of light becomes important at these scales, and for a very large circuit, it may impede the parallel performance because of limits on how closely the circuits can be packed, and of course the power consumption, fan in and fan out, and heat dissipation come into play at high density and high speed high connectivity circuits.

The same concept applies to all sorts of other things. For example, access control. In access control, we may have some number of subjects (S) being controlled with respect to some number of objects (O) with respect to some number of authorizable activities (A). Completely specifying such access control in general takes M*N*A bits, each representing a 3-tuple of subject, object, and activity. My desktop computer, for example, has about 4 million files, each with let's just say 10 bits associated with different activities (read, write, execute, delete, append, etc.) and a few users. So for each user, there are 40 million protection bits that need to be defined to fully control access to files. Consider Google, which processes trillions (10^12) of emails and has billions of users (10^9) using more than 100 different services. At this level of granularity, that's more than 10^23 bits of access control. Today it is estimated that there are perhaps 10 zetabytes (10^22) of data in the world, and at 8 bits per byte, that comes to as many bits for access control at Google as there is total data in the World. And that would just be enough to make precise access control decisions for one large enterprise at the level of access of people to emails with applications.

This level of control is not even a very high level of granularity. Many databases control access at the granularity of records or fields within records, and many databases today have far more than trillions of records. The maximum granularity of control in the digital realm is at the level of the individual digital gate (AND / OR / NOT), since below that level, the circuits are analog. A quad core processor today has perhaps 500M gates.

Granularity of control and adding dimensions lead to the modeling problem

There are bigger problems with more complex authorization schemes. Modern systems seek to use multiple factors for authentication, which then translates into more and more bits required to precisely represent what is allowed under what set of circumstances. Time, for example, can be used to determine accessibility. Time of day has long been used, but consider time and distance requirements for access control (published some years ago as part of a DARPA research project), or perhaps granting access to one record in a database only during a short time window after a request was made from a user through a Web server. Any number of such schemes could be used, but of course then we have the time granularity problem, and in that dimension, computers can discern time at 10^12 or more steps per second. And don't forget that the speed of light comes into play as light can travel only 3*10^8 meters per second, or about a foot per nanosecond. Again, granularity has to be limited for complexity reasons. GPS location has lots of bits as well, and again, at distance intervals of only a foot or so (the approximate depth of a human body), there are lots and lots of places people could be over time.

At any level of granularity we choose, there are lower levels of granularity that we do not control, and of course that means that within the granularity of choice, there are a myriad of things that are being authorized over which we have no specific or differentiated control. Since "maximum granularity" in multiple dimensions for digital systems is so vast, and since within any non-maximum granularity system there are deviations that are not controlled, we can neither attain maximum granularity control nor be perfect in any such controls.

Further, at any level of granularity you choose, you need to control access to the access control bits, and of course then control the bits that control of access control, and so forth in an infinite regress. The same will apply to all sorts of other control attempts with similar characteristics. As such, complexity, granularity, time, space, power, heat, and dimensionality are always considerations in any actual implementation to be considered.

Granularity, Complexity, Resources, and Modeling

For these reasons, we conclude that imperfect modeling is necessary, and of course all actual systems use imperfect models to operate at some level or their scope is limited by the complexity of the modeling. This applies, for example, to simulation, big data analysis, and any predictive mechanism. If we model something small enough we can increase the granularity to a degree that complexity does not overwhelm available resources and get closer and closer to perfection, never actually reaching it because time and space are, according to the current theory of physics, of unlimited granularity.

The translation between management and technology

Which brings us to management of all this stuff. Even if we could find some magical way to allow for all of this control at high granularity (or even lower granularity), there is a completely different problem we face. Whatever the bits are that determine what is allowed, the purpose of all those bits and all that control is to have a system that does what we desire it to do. In the business context, as in the personal context, the desire is expressed by people whose expressions have to be translated into those bits in order for the desires to be reflected in the actions of the systems. People cannot make that many decisions. And even if they could, they cannot get them right. People are not perfect, and digital systems are brittle.

Some translation problems

Digital systems are sequential machines. As such, they perform sequences of combinational operations (those 2-level circuits for example) storing intermediate results and then going on to the next step in the sequence, and so forth. Because of the physics of digital information, any bit flip (which happens from time to time by accident as a result of the underlying analog physics of digital systems) can produce a change in sequence that directs the machine to do arbitrary operations completely different from those without the bit flip. I call this brittle, in the sense that analog systems usually reflect a small input change in a small output change (smoothness), and even in modalities with positive feedback, the sequences of events evolve smoothly; while digital systems compose these changes into bimodal (for binary digital) conditions on each bit at finite time granularity. The brittle nature of digital systems means that a seemingly trivial mistake often results in catastrophic changes with a very wide range of potential outcomes. Analog systems, while subject to such brittle events from time to time, normally produce small changes in outcomes from small changes in decisions-making. So if you start to slip down that slope, you usually compensate in time to not fall over the cliff (unless you want to fall over it, in which case you probably find a way to do so).

People communicate with words, actions, expressions, etc. Computers communicate with bit sequences. Because the bit sequences are precise to the level of granularity of the bit, exact copies can be made and precise meanings, in terms of sequential machine states and outputs, transferred and invoked. Because human communication is not direct brain to brain, and perhaps because even if it were, your brain and mine are not identical in the way we model things, we simply cannot fully express ourselves to each other to communicate our desires, nor do we precisely know what our desires are for the most part. The translation from desires to communication being imperfect, the media through which we communicate being imperfect, our cognitive mechanisms for interpretation being imperfect, and language not expressing ideally what we are thinking, all combine with other such imperfections to produce an inability to express our desires, and perhaps even to understand what they are. Trying to translate this into the digital mechanisms of computers so that they can make moment to moment decisions at a rate higher than we can execute and believing that we can somehow do this to a level we will ultimately agree with is good fiction, but bad reality.

So we use models that we think we understand, and write programs that we think the computer understands, to implement those models in time frames we cannot fathom, and make decisions that we should by now know cannot always be made meeting our desires. We know all of this is imperfect and it will always be so, and as such, we should recognize the fools errand we have beset ourselves upon by ceding so much control over so many things to computers, in the context of humanity that is itself deeply flawed in at least that we fight and kill each other in a race to mutual exploitation and destruction.

People we disagree with know all these things and take advantage of them to take advantage of us. And we use these things to take advantage of them. Any notion that we can design systems to solve these problems without solving the deeper problems within ourselves is problematic and seemingly certain to be problematic. Any notion that somehow the computers will solve this problem for us should be readily overcome by the early attempts to use AI to converse by learning from people that produced rampantly racist remarks and attributes within hours of first coming online.

Paths forward?

“The problem is not in our stars but in ourselves.”

Having gone from technical issues in computers to human frailty, and having made rampant assumptions that current (analog) physics is likely pretty close to right, you may think me a pessimist. But I have some suggestions of paths forward that may help to realize just how stuck we are.

I recently published the last of my series of articles on how to defeat your [whatever][6], I quote:

> **How to defeat any system**
>
> Step 1: Identify the assumptions
>
> Step 2: Violate them
>
> **Conclusions**
>
> You lose.

And I think this may be the way around all of these problems. I think we have long made assumptions that are deeply embedded in our thinking and that we have to break out of if we are to move forward in the new era. I am going to outline some of them here as a starting point with the realization that even if they are bad assumptions we can break free of, it is likely to take longer than I have left in my life or you in yours to see them realized. Perhaps it starts with science fiction...

Some possibly bad assumptions and paths forward:

- The analog world is infinite in granularity in space / time / or any other way
  - What if modern physics is wrong about the smoothness and infinite granularity assumptions it makes? Suppose we found an actual limit to granularity? Of course this would be impossible to prove as it might merely be a limitation of our instrumentation, but still… what if?
    - We might reunify physics around the physics of digital information and actually get to the limits of maximum granularity in everything.
- Digital systems can effectively model human desires
  - What if we decided to go (back) to analog systems for at least come of the things we ask automation to do. What is the physical manifestations of neurons and other body parts built by the reality surrounding us encodes more than just the signals we observe going though dendrites and the chemical levels in the brain, but also encodes the analog inputs we received through out analog neural systems, and the shapes and paths built as we learned encode far more than we have previously understood?
    - We might be able to create models that actually did what humans do in terms of thinking, and perhaps in terms of replacement parts and fusion of feelings and improved communications and so forth. Of course then computers might be as flawed (or is it perfected) as we are.

6 Look for 2021-03C - How to defeat any system on all.net http://all.net/Analyst/2021-03C.pdf

- People are and always will be [name your deadly sins / flaws / etc]
  - What if we had a system of education and communication that built better people and those people decided to work together for mutual benefit instead of filling themselves with hatred and the desire to take advantage of others? What if we came to a different philosophy altogether? What if we decided all those sins and flaws were no such thing? What if we simply decided to turn it into a free for all?
    - We might find a new philosophy (or a bunch of them) that would make our lives better (or worse) and get around all of the problems we have with computers, perhaps specifically the way we seem to care about them and put ourselves into them.
- Some sort of stable hierarchy is necessary for our well being
  - What if we decided to make everything about popularity? What if instead of government and governance, we simply said form relationships at your own risk, caveat emptor. What if we decided that if you get mad at people, you can just go ahead and kill them or whatever? Why not just go to tribalism? It has worked in Afghanistan for centuries.
    - Perhaps a networked approach that gains some of the advantages of hierarchy for reduction of complexity could work, and perhaps we could do a lot better than the representative democracy notion that we have built. Perhaps direct democracy with rapid ability to change as the vote changes, or perhaps with some sort of hysteresis effect on changes to retain a level of stability? Perhaps we could apply this to our systems and vote on who can do what when?
- The minutia is important
  - What if we decided to make the minutia unimportant in the overall context?
    - Instead of getting all the bits right, what if we started to architect and design things so that the consequences of small changes was always small?
- The full faith and credit of stable governments are a sound (our best) basis for trust
  - Of course loss of faith leads to disillusionment leads to susceptibility to influence leads to … the dark side. What if we went to some other basis for trust?
    - Lots of other bases for trust have been tried (and lots have failed as have those trusts in governments). Lots of studies have been done of trustworthiness of people in different contexts, and turning behaviors are often detectable in advance. Perhaps better and more active trust models would be beneficial (how much trust for what based on what, certainly not the misnomer of zero trust or maximum granularity).

Conclusions

This was one of my annual thought pieces to try to break out of the ruts we are (I am) in. It's a good idea to rethink things from scratch and to recall the problems we face at a deeper level from time to time. Have fun with it. Now… back to work.

