BREAKING DEFENSE

E-BOOK:

Artificial Intelligence

The frontline of a new age in defense

FROM OUR UNDERWRITER

Raytheon

Table of Contents

BREAKING DEFENSE

How AI Could Change The Art Of War

What happens when Artificial Intelligence produces a war strategy too complex for human brains to understand? Do you trust the computer to guide your moves, like a traveler blindly following GPS? Or do you reject the plan and, with it, the potential for a strategy so smart it's literally superhuman?

The Pentagon wants AI to assist human combatants, not replace them. The issue is what happens once humans start taking military advice — or even orders — from machines.

The inset image shows what the soldier can see through the wirelessly linked ENVG-III goggle and FWS-I gunsight.

The reality is this happens already, to some extent. Every time someone looks at a radar or sonar display, for example, they're counting on complicated software to correctly interpret a host of signals no human can see. The Aegis air and missile defense system on dozens of Navy warships recommends which targets to shoot down with which weapons, and if the human operators are overwhelmed, they can put Aegis on automatic and let it fire the interceptors itself. This mode is meant to stop massive salvos of incoming missiles but it could also shoot down manned aircraft.

Now, Aegis isn't artificial intelligence. It rigidly executes pre-written algorithms, without machine learning's ability to improve itself. But it is a long-standing example of the kind of complex automation that is going to become more common as technology improves.

While the US military won't let a computer pull the trigger, it is developing target-recognition AI to go on everything from recon drones to tank gun sights to infantry goggles. The armed services are exploring predictive maintenance algorithms that warn mechanics to fix failing components before mere human senses can detect that something's wrong, cognitive electronic warfare systems that figure out the best way to jam enemy radar, airspace management systems that converge strike fighters, helicopters, and artillery shells on the same target without fratricidal collisions. Future "decision aids" might automate staff work, turning a commander's general plan of attack into detailed timetables of which combat units and supply convoys have to move where, when. And since these systems, unlike Aegis, do use machine learning, they can learn from experience — which means they continually rewrite their own programming in ways no human mind can follow.

Sure, a well-programmed AI can print a mathematical proof that shows, with impeccable logic, how its proposed solution is the best, assuming the information you gave it is correct, one expert told the War College conference. But no human being, not even the AI's own programmers, possess the math skills, mental focus, or sheer stamina to double-check hundreds of pages of complex equations. "The proof that there's nothing better is a huge search tree that's so big that no human can look through it," the expert said.

Developing explainable AI — artificial intelligence that lays out its reasoning in terms human users can understand — is a high-priority DARPA project. The Intelligence Community has already had some success in developing analytical software that human analysts can comprehend. But that does rule out a lot of cutting-edge machine learning techniques.

Weirder Than Squid

Here's the rub: The whole point of AI is to think of things we humans can't. Asking AI to restrict its reasoning to what we can understand is a bit like asking Einstein to prove the theory of relativity using only addition, subtraction and a box of crayons. Even if the AI isn't necessarily smarter than us — by whatever measurement of "smart" we use — it's definitely different from us, whether it thinks with magnetic charges on silicon chips or some quantum effect and we think with neurochemical flows between nerve cells. The brains of (for example) humans, squid, and spiders are all more similar to each other than either is to an AI.

Alien minds produce alien solutions. Amazon, for example, organizes its warehouses according to the principle of "random stow." While humans would put paper towels on one aisle, ketchup on another, and laptop computers on a third, Amazon's algorithms instruct the human workers to put incoming deliveries on whatever empty shelf space is nearby: here, towels next to ketchup next to laptops; there, more ketchup, two copies of 50 Shades of Grey, and children's toys. As each customer's order comes in, the computer calculates the most efficient route through the warehouse to pick up that specific combination of items. No human mind could keep track of the different items scattered randomly about the shelves, but the computer can, and it tells the humans where to go. Counterintuitive as it is, random stow actually saves Amazon time and money compared to a warehousing scheme a human could understand.

In fact, AI frequently comes up with effective strategies that no human would conceive of and, in many cases, that no human could execute. Deep Blue beat Garry Kasparov at chess with moves so unexpected he initially accused it of cheating by getting advice from another grandmaster. (No cheating — it was all the algorithm). AlphaGo beat Lee Sedol with a move that surprised not only him but every Go master watching. Libratus beat poker champions not only by out-bluffing them, but by using strategies long decried by poker pros — such as betting wildly varying amounts from game to game or "limping" along with bare-minimum bets — that humans later tried to imitate but often couldn't pull off.

If you reject an AI's plans because you can't understand them, you're ruling out a host of potential strategies that, while deeply weird, might work. That means you're likely to be outmaneuvered by an opponent who does trust his AI and its "crazy enough to work" ideas.

At what point do you give up on trying to understand the alien mind of the AI and just "hit the I-believe button"?

The New Principles of War

If you do let the AI take the lead, several conference participants argued, you need to redefine or even abandon some of the traditional "principles of war" taught in military academies. Now, those principles are really rules of thumb, not a strict checklist for military planners or mathematically provable truths, and different countries use different lists. But they do boil down centuries of experience: mass your forces at the decisive point, surprise the enemy when possible, aim for a single and clearly defined objective, keep plans simple to survive miscommunication and the chaos of battle, have a single commander for all forces in the operation, and so on.

To start with, the principle of simplicity starts to fade if you're letting your AI make plans too complex for you to comprehend. As long as there are human soldiers on the battlefield, the specific orders the AI gives them have to be simple enough to understand — go here, dig in, shoot that — even if the overall plan is not. But robotic soldiers, including aerial drones and unmanned warships, can remember and execute complex orders without error, so the more machines that fight, the more simplicity becomes obsolete.

The principle of the objective mutates too, for much the same reason. Getting a group of humans to work together requires a single, clear vision of victory they all can understand. Algorithms, however, optimize complex utility functions. For example, how many enemies can we kill while minimizing friendly casualties and civilian casualties andcollateral damage to infrastructure? If you trust the AI enough, then the human role becomes to input the criteria — how many American soldiers' deaths, exactly, would you accept to save 100 civilian lives? — and then follow the computer's plan to get the optimal outcome.

Finally, and perhaps most painfully for military professionals, what becomes of the hallowed principle of unity of command? Even if a single human being has the final authority to approve or disapprove the plans the AI proposes, is that officer really in command if he isn't capable of understanding those plans? Is the AI in charge? Or the people who set the variables in its utility function? Or the people who programmed it in the first place?

GENOCIDE SWARMS & ASSASSIN DRONES:

The Case For Banning Lethal AI

70-ton robotic battle tanks? Scary. Three grams of explosive on a mini-drone that knows your face? Also scary. Thousands of such drones? Millions? That's potentially a strategic game-changer in a way that automating conventional military hardware is not.

"I'm not too worried about vast autonomous swarms of battle tanks," said Berkeley AI scientist and activist Stuart Russell.

This was surprising considering Russell had criticized the US Army's ATLAS project to put Artificial Intelligence in armored vehicles, a system intended to assist human gunners that he argued could all too easily replace them altogether. Quartz.com headlined its story on ATLAS "The US Army wants to turn tanks into AI-powered killing machines." Okay, so the US Army actually doesn't want that at all — replacing loyal, well-trained soldiers with unproven technology justifiably gives generals the heebie-jeebies — but just the possibility of robot tanks got a lot of people pretty worried.

A quadcopter that slipped through security to land on the White House lawn

Russell, however, has bigger things to worry about — or rather, much, much smaller things.

"I think of autonomous tanks as mainly a weapon for war between major powers," he said. Taking the humans out of an armored vehicle, fighter jet, or warship could make it more effective in combat, and, because you no longer need space and life-support for human crew, it can definitely make them smaller and cheaper. But automating conventional war machines doesn't make them smaller and cheaper enough that governments can stockpile vast swarms of them in secret and smuggle them into an enemy capital, or that terrorists can build them in garages with 3D printers.

So what Russell really worries about is not robotic tanks — though he'd definitely prefer a world without them — but what happens when the technology is developed and the precedent is set.

"Given the cost of a new M1A2 around $9 million...there are far cheaper ways to flatten a city and/or kill all of its inhabitants," Russell told me. "The problem with full autonomy is that it creates cheap, scalable weapons of mass destruction."

It's already possible to build assassin drones by combining off-the-shelf quadcopters, small amounts of homebrewed explosive, and the kind of facial-recognition technology Facebook uses to tag other people's bad pictures of you.

"My UAV colleagues tell me they could build a weapon that could go into a building, find an individual, and kill them as a class project," Russell said. "Skydio plus self-driving cars plus AlphaStar more or less covers it." (Skydio's a drone you can buy on Amazon; AlphaStar is a version of the DeepMind AI that beats humans at complex strategy games like Starcraft). In fact, he said, Switzerland's domestic security agency, DDPS, "made some to see if they would work — and they do."

Not only would they work, they've already been tried. ISIS has already used mini-drones as "flying IEDs," and someone attempted to assassinate Venezuelan president Nicolàs Maduro with a pair of exploding drones.

Soldier with handheld quadcopter

Small Drones, Big Kills

Now what happens when you scale this up? Russell and fellow activists actually produced a video, Slaughterbots, in which swarms of mini-drones attack, among other groups, every member of Congress from a particular party. But that's still thinking small.

Remember, once you've written the software, you can make infinite copies; lone cranks can make explosives; and mini-drones are getting cheaper by the day. Remember also that the Chinese government has personal information on some 22.1 million federal employees, contractors, and their family members from the Office of Personnel Management breach two years ago. Now imagine one out of every thousand shipping containers imported from China is actually full of mini-drones programmed to go to those addresses and explode in the face of the first person to leave the house. Imagine they do this the day before China invades Taiwan. How effectively would the US government react?

A rogue state or terrorist group could go further. How about programming your mini-drones to kill everyone who looks white, or black or Asian? (One Google facial recognition algorithm classified African-Americans as "gorillas," not humans, so racist AI is a mature technology). It would be genocide by swarm.

Such a tactic might only work once, much like hijacking airliners with box cutters on 9/11. "Small drones are vulnerable to jamming, to high-powered microwaves, to other drones that might intercept them, to nets," said Paul Scharre, an Army Ranger turned thinktank analyst. "Bullets work pretty well… I have a buddy who shot a drone out of the sky back in Iraq in 2005." (Unfortunately, the drone was American). At least some object-recognition algorithms can be tricked by carefully applied reflective tape.

"People are working on countermeasures today," Scharre told me, "and the bigger the threat becomes, the more people have an incentive to invest in countermeasures."

But how do you stop tiny drones from becoming a big threat in the first place? While technology to build a "working prototype" already exists, Russell told me, the barrier is mass production.

No national spy agency or international monitoring regime can find and stop everyone trying to make small numbers of drones. But, Russell argues fervently, a treaty banning "lethal autonomous weapons systems" would prevent countries and companies from openly producing swarms of them, and a robust inspection mechanism — perhaps modeled on the Organisation for the Prohibition of Chemical Weapons — could detect covert attempts at mass production.

A quadcopter drone destroyed by Rafael's "Drone Dome" laser system

Without a ban, Russell said, legal mass production could make lethal swarms as easy to obtain as, say, assault rifles — except, of course, one person can't aim and fire thousands of rifles at once. Thousands of drones? Sure.

So don't fear robots who rebel against their human masters. Fear robots in the hands of the wrong human.

Would a ban on lethal AI actually work? Would the United States actually want it to work?

ARTIFICIAL INTELLIGENCE:
Are We Losing The Race?

A few hours before the Pentagon released its first Artificial Intelligence strategy in February 2019, SASC Chair James Inhofe was asked why the US military — and the US generally — appeared to be doing so relatively little about it, while China has made AI the centerpiece of an outright societal realignment, complete with a master plan and huge amounts of targeted money.

"I think Russia and China are in a better position than we are at the moment on Artificial Intelligence," the senator said straightaway.

He was asked if he would press hard for more money.

Answer: "To me, there are other things that need to be done first," Sen. Inhofe said. As the senator from Oklahoma, he pointed to one of the subjects he knows best: artillery — or, in the current jargon, Long-Range Precision Fires, the US Army's No. 1 modernization priority. (Of course, Fort Sill is in the senator's state, home to the Army's Field Artillery and Air Defense Artillery schools.) He pointed to recent statements by Gen. Joseph Dunford, Chairman of the Joint Chiefs, that we have lost our quantitative and quantitative edges in artillery, and by Gen. Mark Milley, Army Chief of Staff and Dunford's probable successor, who said the US is outranged and outgunned by our adversaries. Inhofe also mentioned the readiness woes that afflict the aging F-18 fleet, among others. So, don't expect a great push for either more money or more focus on AI from the SASC chairman.

When former Deputy Defense Secretary Bob Work heard the then head of Google's parent company, Eric Schmidt, say in November 2017 that America needs a national strategy for developing Artificial Intelligence, one image sprang to his mind's eye.

"The image that popped into my mind was of Nikita Khrushchev banging his shoe in the UN and saying, 'We will bury you," Work said. "As Eric said, the US does not have a coherent strategy" for developing AI, the father of the Pentagon's Third Offset Strategy opined.

Well, the Pentagon has released its AI strategy. There are virtually no mentions of increased funding in it. But it does make the stakes clear: "Failure to adopt AI will result in legacy systems irrelevant to the defense of our people, eroding cohesion among allies and partners, reduced access to markets that will contribute to a decline in our prosperity and standard of living, and growing challenges to societies that have been built upon individual freedoms."

The Pentagon

It identifies the center of gravity for AI work in the military, the Joint Artificial Intelligence Center (JAIC). And it sets some priorities:

"We will launch a set of initiatives to incorporate AI rapidly, iteratively, and responsibly to enhance military decision-making and operations across key mission areas. Examples include improving situational awareness and decision-making, increasing the safety of operating equipment, implementing predictive maintenance and supply, and streamlining business processes. We will prioritize the fielding of AI systems that augment the capabilities of our personnel by offloading tedious cognitive or physical tasks and introducing new ways of working."

Importantly, the strategy notes the importance of ethics in developing and using AI, saying the Pentagon "will articulate its vision and guiding principles for using AI in a lawful and ethical manner to promote our values. We will consult with leaders from across academia, private industry, and the international community to advance AI ethics and safety in the military context."

Then it sort of issues a laundry list of what AI will be used for, promising to "share our aims, ethical guidelines, and safety procedures to encourage responsible AI development and use by other nations."

- Increasing safety of operating equipment. AI also has the potential to enhance the safety of operating aircraft, ships, and vehicles in complex, rapidly changing situations by alerting operators to hidden dangers.

- Implementing predictive maintenance and supply. We will use AI to predict the failure of critical parts, automate diagnostics, and plan maintenance based on data and equipment condition. Similar technology will be used to guide provisioning of spare parts and optimize inventory levels. These advances will ensure appropriate inventory levels, assist in troubleshooting, and enable more rapidly deployable and adaptable forces at reduced cost.

- Streamlining business processes. AI will be used with the objective of reducing the time spent on highly manual, repetitive, and frequent tasks. By enabling humans to supervise automated tasks, AI has the potential to reduce the number and costs of mistakes, increase throughput and agility, and promote the allocation of DoD resources to higher-value activities and emerging mission priorities.

And they will work on solving hard "global challenges of significant societal importance" such as how to US AI for humanitarian assistance and disaster relief for wildfires, hurricanes, and earthquakes. "These open missions will challenge a broad community to advance the state of AI and learn how to operationalize the technologies on an integrated basis across domestic and international organizations. They will contribute to the development of thousands of new AI experts needed for public service over the next decade and spur future AI progress across multiple sectors," the strategy promises.

In addition to the Pentagon strategy, the White House announced an Executive Order yesterday designed to coordinate AI work across the federal government. It contained few details, in contrast with the Chinese plan. One of the few former defense policymakers with real world experience in AI, Wendy Anderson, gave the Trump Administration credit for the EO, saying it "is significant." But: "That said, the content of the EO isn't new. <any AI thought leaders, technologists, others from the tech community, scientists, and policymakers have been making these important points for years. If there's no implementation plan behind the EO – with details, deadlines, and funding — then it may be worse than no EO at all."

Anderson, who was deputy chief of staff for Defense Secretary Chuck Hagel, put her finger on the one thing that is, at least in public, missing from all this: "I'd like to see the implementation plan and resources behind it as soon as possible. If we don't want to fall behind on this game-changing technology, we need to up our game, and we need to do so now," she says in her email. "In contrast to our EO released today, the Chinese strategy, which is full of details, deadlines, and a clear funding plan, also has engaged support and action at the very top."

Anderson, who now works for an Austin,Texas company called SparkCognition that designs and builds AIs, asks what has the US been doing compared to the Chinese? "To date, we have mostly engaged in debates about banning AI exports. In the absence of significant US governmental AI spend and in the absence of a robustly resourced national AI strategy, we are now also attempting to limit our private companies' ability to access capital via international sales to the world's largest markets."

Bottom line, for Anderson? "We are losing the money/investment race big time."

Readers, are we missing our second Sputnik moment?

FEAR & LOATHING OF AI:

How The Army Triggered A Fear Of Killer Robots

The April 2018 Army released its ATLAS artificial intelligence targeting program solicitation, inspiring runaway headlines about "AI-powered killing machines." Why did this happen? The answer lies in a strange mix of misperceptions and some very real loopholes in the Pentagon's policy on lethal AI.

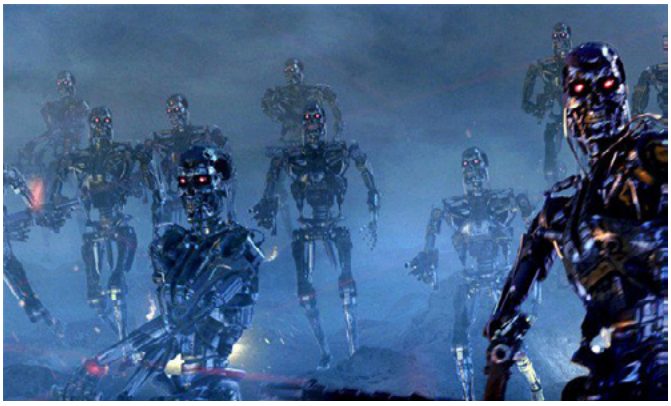

No, the US Army is not building this

"The US Defense Department policy on autonomy in weapons doesn't say that the DoD has to keep the human in the loop," Army Ranger turned technologist Paul Scharre said. "It doesn't say that. That's a common misconception."

Buzzwords & Firestorms

ATLAS came to public attention in about the worst way possible: an unheralded announcement on a federal contracting website (fbo.gov) on February 19th 2019, an indigestible bolus of buzzwords that meant one thing to insiders but something very different to everyone else — not just the general public but even civilian experts in AI.

The name itself is ominous: ATLAS stands for Advanced Targeting and Lethality Automated System. The wording on the website made it worse, soliciting white papers on "autonomous target acquisition technology, that will be integrated with fire control technology, aimed at providing ground combat vehicles with the capability to acquire, identify, and engage targets at least 3X faster than the current manual process."

"The LA in ATLAS stands for Lethality Automated," pointed out an appalled Stuart Russell, an AI scientist at Berkeley who's campaigned for a global ban on lethal autonomous weapons. "'Acquire, identify, and engage targets' is essentially the UN definition of lethal autonomy."

But it's not the military definition, which is where the problem starts.

The military has long applied the loaded word "lethality" to anything that could make weapons more effective, not just the weapons themselves. Adding new infrared targeting sensors to tanks, for example, is officially a "lethality" upgrade. Networking Navy ships so they can share targeting data is called "distributed lethality." Then came Defense Secretary Jim Mattis, a retired Marine Corps four-star who liked the word "lethal" so much that underlings plastered it on everything they were trying to sell him on, from high-tech weapons to new training techniques.

What about "engagement"? In plain English, a "military engagement" means people are trying to kill each other (lethally)." But in the military, "engagement" can mean anything from "destroy" to "consider" to "talk to." A Key Leader Engagement (KLE) in Iraq meant soldiers talking with a tribal elder, sheikh, or other influential person over tea.

So in military language — at once abstrusely technical and sloppy — an artificial intelligence can increase "lethality" and "engage" a potential target by helping a human soldier spot it and aim at it, without the AI having any control over the trigger.

There are people in the Pentagon, however, who were aware of how this all sounded even before the original "killing machines" story came out on Quartz. In fact, within hours of the ATLAS solicitation going online, the head of the Pentagon's nine-month-old Joint Artificial Intelligence

Center, Air Force Lt. Gen. John Shanahan, was contacting Army counterparts trying to head off what he feared would be a "firestorm" of negative news coverage.

As far as I can determine, the Army hadn't officially informed JAIC about the relatively small and nascent ATLAS project. Instead, someone — we don't know who, but they weren't on the JAIC staff itself — spotted the online announcement almost immediately and raised a red flag. That JAIC not only got that information but actually acted on it so quickly is a remarkable feat for any government agency, let alone one created less than nine months ago: Where the usual bureaucratic channels dropped the ball, JAIC picked it up.

Unfortunately for the Pentagon, JAIC and the Army didn't move fast enough to get Quartz to update its story. Instead, the next story was in Defense One, headlined "US Military Changing 'Killing Machine' Robo-tank Program After Controversy." In fact, as the body of the article explained, the change was to the wording of the solicitation, not to what the program was actually doing.

The revised solicitation for ATLAS adds a paragraph emphasizing the system will be "consistent with DoD legal and ethical standards," especially Department of Defense Instruction 3000.09 on "Autonomy in Weapon Systems." The final decision to fire will always be a human being's job, the Army insists, in keeping with Pentagon policy.

But policy is not law, and the Pentagon leadership can change it unilaterally. What's more, even though the military's AI policy is usually described as requiring a "human in the loop," there's actually an enormous loophole.

"It authorizes the development of weapons that use autonomy...for defensive purposes like in Aegis or Active Protection Systems," Scharre said. "For anything else, it creates a review process for senior leaders to make a determination."

"It's not a red light," Scharre told me. It's a stop sign: You halt, you check out the situation — and then you can go.

The Problem With Policy

Are you worried the US military will give computers control of lethal firepower? Well, in one sense, you're too late — by decades. Scores of Navy warships use the Aegis fire control system to track and target potential threats in the air. Normally a human has to press the button to fire, but the sailors can also set the computer to launch interceptor missiles on its own. That's an emergency option, intended for use only when the human crew can't keep up with massive salvos of incoming missiles — but it could shoot down manned aircraft as well.

Aegis isn't the only example, Scharre pointed out. There is the Navy's Phalanx and its Army spin-off, C-RAM, which automatically shoot down incoming missiles and rockets. The Army's started fielding Active Protection Systems, a miniaturized missile defense that can fit on a tank.

None of these systems is an artificial intelligence in the modern sense. They are purely deterministic sets of old-fashioned algorithms that always produce the same output from a given input, whereas machine learning algorithms evolve — often unpredictably and sometimes disastrously — as they process more and more data. The initial version of Aegis was actually introduced in 1973, long before the Defense Department first issued DoD Instruction 3000.09 in 2012. But it's not only old systems being grandfathered in: Active Protection Systems are just entering service now.

So what does the regulation actually say?

- DoD 3000.09, Section 4.c(2), covers "human-supervised autonomous weapons systems" — since a human overseer can turn it off at any time , like Aegis– and specifically limits them to defensive purposes, explicitly banning the "selecting of humans as targets."

- Section 4.c(3) allows computer-controlled non-lethal systems, such as radar jammers. (Automated cybersecurity software is permitted elsewhere).

- Section 4.c(1) allows the use of lethal force by "semi-autonomous weapons systems" (emphasis added), which aren't fully computer-controlled. But even those must "not autonomously select and engage individual targets or specific target groups that have not been previously selected by an authorized human operator."

Such strictly regulated systems are a far cry from the Terminator, or even Stuart Russell's more realistic nightmare scenario of swarming mini-drones. But while Section 4.c is the heart of the Pentagon policy on autonomous weapons, it's immediately followed by a loophole:

- Section 4.d states that "Autonomous or semi-autonomous weapon systems intended to be used in a

manner that falls outside the policies in subparagraphs 4.c.(1) through 4.c.(3) must be approved" before development can proceed. Who approves? Two deputy secretaries of defense (policy and technology) and the Chairman of the Joint Chiefs. Getting three such high-level officials to sign on is a daunting challenge for any bureaucrat, but it's hardly impossible.

- Even after the three officials approve an exception, the system must follow a long list of safety and testing guidelines and ensure "commanders and operators [can] exercise appropriate levels of human judgment in the use of force." But "appropriate" is left undefined.

A soldier holds a PD-100 mini-drone during the PACMAN-I experiment in Hawaii.

What's more, if all three officials agree, they can ask the Deputy Secretary of Defense to waive all of those restrictions, "with the exception of the requirement for a legal review, in cases of urgent military operational need" — again, left undefined.

Nowhere in this document, incidentally, will you find the comforting but imprecise phrase "human in the loop." In fact, when I used it in a query to the Pentagon, I got gentle chiding from DoD spokesperson Elissa Smith: "The Directive does not use the phrase 'human in the loop,' so we recommend not indicating that DoD has established requirements using that term."

The Real Barrier

So what is stopping the Defense Department from developing AI weapons that can kill humans? The real barrier, it turns out, is not legal or technological: It's cultural. The US military isn't developing killer robots because it doesn't want them.

Every officer and official I've ever talked to on the subject, for at least eight years, has said they want AI and robotics to help the human, not replace them — and even then, they want AI primarily in non-combat functions like logistics and maintenance. In fact, Pentagon leaders seem to think taking the human out of the loop would be giving up one of American military's most crucial advantages: the training, creativity, and, yes, ethics of its people.

"The last thing I want is you to go away from this thinking this is all about technology," then-Deputy Secretary Robert Work told us in 2015. Work, whose Third Offset Strategy first made AI a top priority for the Pentagon, has remained deeply engaged in the debate. "The number one advantage we have is the people in uniform, in our civilian work force, in our defense industrial base, and the contractors who support us."

But Work also said "we want our adversaries to wonder what's behind the black curtain," Stuart Russell pointed out, as part of a deterrence strategy. Does the waiver provision in Pentagon policy means those secret programs could already include lethal AI?

Well, no, Smith told me in a statement: "To date, no weapon has been required to undergo the Senior Review in accordance with DOD Directive 3000.09."

But with the stakes so high, Russell argues, can the Pentagon really expect potential adversaries or even US-based companies like Google to take it at its word? Or, following the longstanding intelligence maxim to look at capabilities instead of intentions, should they judge programs like ATLAS, not by what the US says they'll do, but by what they could become?

"The declared intention and the intention are not the same thing," Russell said bluntly. "If the aggressive pursuit of partial autonomy were accompanied by a full-on diplomatic effort to negotiate an international ban on full autonomy, there would be less of an issue. As it stands, the ATLAS announcement will be taken as an indicator of future intent."

Artificial Intelligence Can Help National Security Agencies Respond to Disasters

AI CAN SAVE LIVES WHEN DISASTER STRIKES

By Christopher A. Worley

According to the World Health Organization more than 160 million people are significantly impacted by natural disaster each year.

When these disasters strike struggling nations, the results could be catastrophic, leading to high casualty rates, disease outbreaks and serious risks to global and regional security. Even in the U.S., a nation well prepared to weather significant disasters, disasters can pose significant readiness challenges if major military installations are impacted.

In both cases, first responders –especially militaries and other national agencies – need to respond quickly, a process that would be easier if they had access to artificial intelligence tools.

We already benefit from the integration of AI and machine learning now – from your car's navigation system to your digital home assistants to Netflix's surprisingly accurate recommendation engine. But, the potential of AI extends far beyond convenience; AI can save lives.

When hurricanes and typhoons strike, entire communities are at risk – not just from the winds, rain and storm surge, but the secondary impacts on power and communication networks. AI can expedite emergency response and help prioritize where first responders are needed most.

Today's technology and algorithms can assess vulnerable power lines based on modeled wind speeds before the storm strikes. This could allow power companies to alert hospitals, retirement communities and cellular network managers, giving them more time to move generators into place and plan evacuations.

AI-enabled drones, equipped with temporary cellular phone relays, could restore critical communication channels quickly. Other drones, equipped with optical sensors, could search for stranded victims, then automatically pass their positions to rescue crews.

When it comes to natural disasters, few move as quickly and unpredictably as forest fires. Advanced AI could image terrain and weather data and develop real-time, dynamic escape routes to those desperately trying to flee.

AI is maturing rapidly and has the potential to benefit our society in ways both big and small. Don't get me wrong, I truly appreciate when my car knows how to avoid traffic jams when there is construction ahead.

But, I will appreciate it even more when a member of the National Guard can save a member of my family following a natural disaster because this technology was able to locate them and speed their rescue.

Now is the time to focus our efforts and resources at equipping first responders with AI-enabled technology. It does more than make their heroic efforts easier, it expands their impact and benefits everyone at risk.

Christopher Worley is Director of Digital Innovation at Raytheon Intelligence, Information and Services

Raytheon

FIX IT BEFORE IT BREAKS:

SOCOM, JAIC Pioneer Predictive Maintenance AI

The Pentagon's Joint Artificial Intelligence Center deployed its first operational project around june 2018, a joint venture with Special Operations Command to predict helicopter breakdowns before they happen. It's considered an example of how SOCOM is applying artificial intelligence to real-world warfare, a cutting-edge role it also played with the intelligence-gathering Project Maven.

Leading the charge since 2016: SOCOM's current commander at the time, Gen. Raymond "Tony" Thomas, who one SOCOM source extolled as "truly the father of AI in SOF" (Special Operations Forces). But in his last Senate hearing before retirement, Thomas himself gave credit to the former CEO of Google for giving him a precision-guided kick in the pants.

Eric Schmidt heads the Defense Innovation Advisory Board meant to infuse outside commercial innovation into the stodgy Defense Department. An immensely experienced alumnus of Google, Alphabet, Sun and Xerox, Schmidt visited SOCOM's Tampa headquarters three years ago, early in Thomas's tour as commander.

"He felt compelled to give me [a] quick assessment," Thomas recalled. "He said you've got tremendous people. You prototype pretty effectively, and you're absolutely terrible — he had some more colorful words than that — for machine learning, applied artificial intelligence."

"Truthfully, it gave me a spark three years ago and turned me into a zealot" for AI, Thomas said of Schmidt's visit. "More importantly, it really has reoriented our command to embrace [artificial intelligence] and apply it."

The working assumption at SOCOM today, Thomas said, is that AI "has relevance to everything we do until it's proven otherwise. So we're taking, not so small bites, but some pretty substantial bites, in embracing, applying artificial intelligence, and I'm excited about where we're going in the future."

Testifying alongside Thomas before the Senate Armed Services Committee was Gen. Paul Nakasone, who runs both Cyber Command and the National Security Agency, someone who'd better be an expert on AI.

"We have already seen the power at the National Security Agency of what artificial intelligence can do for our foreign intelligence mission, our cybersecurity missions," Nakasone said when asked about AI. But rather than give specifics on his own highly classified projects, Nakasone passed the mike to Thomas with a compliment: "Special Operations Command really has led a lot of work in artificial intelligence."

160th Special Operations Aviation Regiment (SOAR) training with Special Boat Team 12

Maven & Maintenance

So what has Thomas done, exactly? Our SOCOM source gave us two examples, although there are probably more deep in the classified world: Project Maven and predictive maintenance.

Project Maven was the brainchild of then-Deputy Defense Secretary Robert Work, who created an Algorithmic Warfare task force to apply AI to urgent military problems. (Ironically, one of the early private-sector partners was Google, which pulled out after protests from almost 4,000 employees). Maven uses AI to analyze intelligence on terrorist suspects, chiefly by combing through drone surveillance footage, and deduce their latest location for drone strikes or Special Operations raids.

ARTIFICIAL INTELLIGENCE:

Will Special Operators Lead The Way?

The Pentagon's new artificial intelligence strategy shows how the military is shifting from old-school heavy-metal hardware – tanks, ships, planes – to a world where software makes the difference between victory and defeat. And the bigger this shift becomes, several experts suggest, the bigger the role for Special Operations Command in pioneering new technology. Then the new Joint Artificial Intelligence Center can cherry-pick the successes and scale them up for wider use.

Sure, SOCOM has a long tradition of innovation in general, but with a $14 billion budget, it can't build aircraft carriers or stealth fighters. (It gets its aircraft from the larger services and modifies them for special missions). What SOCOM can test-drive for the services is the smaller stuff, from off-road vehicles to mini-drones to frontline wireless networks – but in the information age, the small stuff is a big deal.

We're not talking killer robots here, but intangible algorithms that help humans make sense of masses of data. (Much of that data, admittedly, is gathered by drones and other unmanned systems, but most are unarmed and even the armed ones can't fire without a human command). What SOCOM and DoD's AI Strategy as a whole are looking for, fundamentally, is AI software that can rapidly process vast amounts of information on everything from threats to targets to logistics, provide recommendations to commanders, and maybe take instant action against split-second threats like hackingand jamming, but leave life-and-death decisions to human beings – who remain, as the strategy says, "our enduring source of strength."

SOCOM Can Lead The Way

"The SOF guys are less risk averse than conventional ground forces, so they're more apt to push the limit," said Bob Work, former deputy secretary of defense and father of the AI-driven Third Offset Strategy. "Their commanders also have embraced AI and autonomous ops.... so I think all the conditions are set for SOF to lead the way in the more direct combat applications of AI and autonomy."

Special Operations missions are particularly demanding in ways that could benefit from artificial intelligence, Work told me. "Global man-hunting will see new types of AI-empowered human-machine combat teaming" to sort through masses of surveillance data. "Operating in the grey zone" – the ambiguous arena of proxy war and deniable cyber attacks – "will require Ai-empowered pattern recognition. [And] I can see SOF pursuing a wide range of AI-empowered robotic systems, for house clearing. HVT [High Value Target] tracking, dynamic breaching, etc."

The leadership of Special Operations Command has started pushing hard on artificial intelligence, said Wendy Anderson. Now with AI firm SparkCognition, she was chief of staff to Work's old boss, technophile Defense Secretary Ash Carter, back when he was DepSecDef and head of acquisition.

"SOCOM is clearly starting to style itself as an AI Command," Anderson told me. "The SOF community is well positioned to lead the way in the digital space, especially with regards to the operationalization and deployment of AI." The foundation, she said, is SOCOM's unique combination of urgent operational needs, relative lack of bureaucracy, special acquisition authorities and institutional culture less afraid of risk than the mainstream military. But more recently and specifically on AI, she said, SOCOM has made "a number of smart, timely, and innovative decisions by senior leaders there, including, perhaps most prominently, the executive decision to bring on board a Chief Data Officer."

Anderson's assessment echoes a self-confident statement by SOCOM's own director for science & technology, Lisa Sanders, when I asked her about this topic at the annual NDIA SOLIC conference last week.

"The digital space... it's absolutely an area that SOCOM can lead the way," Sanders said. "We have that unique relationship with our Chief Information Officer and Chief Data Officer. SOCOM has our own network: We have the fourth largest network in the Department of Defense" – large enough to be a real test of new technology, small enough to be nimble. Equally important, SOCOM also has the authority to rapidly approve new technologies for operational use, without waiting for a mother-may-I from a service or the Office of the Secretary of Defense . "If the opportunity is worth taking a risk," Sanders said, "we can certify it for use on our network – and we will certify it."

Special Operations command center

It also sounds like key officials in DoD are comfortable with SOCOM taking those risks. "The cyber domain demands a response faster than our traditional models work, there's no doubt about that," said Brig. Gen. Dennis Crall, who's the Pentagon's senior uniformed advisor on cyber policy and joined at the hip with the increasingly powerful DoD CIO, Dana Deasy. "I realize the scale's a little bit different, but I look at how SOCOM... can do rapid prototyping, fielding," Crall said at last week's conference. "They can test very quickly and determine what's right for the warfighter."

All that said, SOCOM can still screw up, cautioned Kara Frederick, who worked as a civilian intelligence analyst at Naval Special Warfare Command — including three deployments to Afghanistan. (She later went on to work for Facebook security before joining the same thinktank, CNAS, where Work now hangs his hat). But when something does go wrong, she said, hardbitten frontline sergeants can get the word back to the generals faster than any other part of the military.

"As the TALOS 'Iron Man' project showed, SOCOM isn't magic when it comes to emerging technologies," Frederick told me, referring to a much-hyped super-suit exoskeleton that the command now admits can't be built any time soon. But SOCOM does have resources, flexibility, and access to top intelligence community talent that conventional forces don't.

Just as important is the institutional culture, she said: "They pride themselves on their 'flat' [organization], which gives the average Special Operator much more agency than your typical E-5 [sergeant] in conventional forces. This means that good ideas are more likely to filter up quickly, and — similarly — leadership will also hear about the bad ideas directly from those actually employing the tech."

Strategy & SOCOM

Now compare these statements about Special Operations Command to what the Defense Department's new artificial intelligence strategy, released just yesterday, says how AI innovation needs to work. The document – at least the 17-page unclassified summary that's been publicly released – never mentions Special Operations by name, but it calls for characteristics that SOCOM shows in spades.

Innovation must come bottom-up, from all over the Defense Department, the strategy says in several places:

"One of the U.S. military's greatest strengths is the innovative character of our forces. It is likely that the most transformative AI-enabled capabilities will arise from experiments at the 'forward edge,' that is, discovered by the users themselves in contexts far removed from centralized offices and laboratories.

"We will encourage rapid experimentation, and an iterative, risk-informed approach to AI implementation.... We are building a culture that welcomes and rewards appropriate risk-taking to push the art of the possible: rapid learning by failing quickly, early, and on a small scale."

"Execution will prioritize dissolving the traditional sharp division between research and operations.....insights must transition immediately to the research venue, and research must benefit by the immediate involvement of end users in the technology development process."

All of this sounds a lot like what SOCOM has been doing for years.

But while SOCOM can blaze trials, it can't pave highways. Scaling up the successful experiments for use across the Defense Department, the strategy says, is the role of the eight-month-old Joint Artificial Intelligence Center:

"**Scaling successful prototypes.** The JAIC will work with the Military Departments and Services and other organizations to scale use cases throughout the Department in a manner that aligns with and leverages enterprise cloud adoption.... The JAIC will strengthen the efforts of the Military Departments and Services and other independent teams across DoD as they continue to develop and execute new AI mission initiatives.... The JAIC will work closely with individual components [of the Defense Department] to help identify, shape, and accelerate their component-specific AI deployments, called 'Component Mission Initiatives' or 'CMIs.'"

Again, the public summary of the strategy never mentions Special Operations by name. But then it doesn't mention Cyber Command either, another organization that strives to be on the cutting edge, albeit with a much shorter history and thus less of a track record than SOCOM. Besides the Joint AI Center itself, in fact, the strategy calls out only two Defense Department organizations by name:

- One is the Defense Innovation Unit, DIU (formerly DIUx for "experimental), which has had real successes bringing Silicon Valley innovation to the armed forces, but is a small outfit not yet three years old. SOCOM has 70,000 people and 31 years of history.

- The other is the Defense Advanced Research Projects Agency, which specializes in high-risk, high-reward research pursuing fundamental breakthroughs that it hands to other agencies to turn into specific weapons. SOCOM, by contrast, doesn't even fund basic research (budget function 6.1) and concentrates on near-term applications, often of technology borrowed directly from the commercial world. So SOCOM has a very different niche, with more potential to make a near-term impact while DARPA works on revolutionizing the future.

SOCOM also has a very close relationship with the Intelligence Community, which often gives it priority access to both technology and data.

Finally, it's worth noting that artificial intelligence will probably be essential to create acommunications network fast, flexible, and robust enough to coordinate far-flung forces operating across the land, sea, air, space, and cyberspace — a concept the Army and Air Force have embraced as Multi-Domain Operations. SOCOM sees itself as well-suited to this new way of warfare, since it already includes elements of all four services operating in all five domains.

"SOCOM is by definition joint and works in multiple domains," Sanders told me after her remarks to the conference. "It's largely a question of scale, because we are in a smaller environment with a specific, focused objective, [but] SOF is actively engaged. In every one of those multi-domain concepts, there will be an element specifically for SOF, and so we have aspects of our command that are responsible for working in those [as] they're developed through exercises and wargames."

The Human Element

For all this proposed change, one thing stays constant: In both the near term and the long, human beings remain central to the American military's approach to artificial intelligence, a hybrid of human and machine sometimes likened to the mythical centaur. The strategy calls for the "thoughtful, responsible, and human-centered adoption of AI in the Department of Defense" (emphasis ours).

Using trust and technology to empower the troops is, of course, another Special Operations tradition.

"The most complex weapon we have on the battlefield is the SOF operator themself," said Cdr. James Clark, a Navy SEAL and now SOF program manager in the Pentagon's Strategic Capabilities Office. And, he told the conference, the only way special operations can succeed, with small teams scattered over vast areas with often erratic communications with each other and their superiors, is trust: "There's trust between the leaders, there's trust in what their SOF operators are capable of doing."

"We're ultimately going to have to develop trust with our machine learning, with our artificial intelligence, and we're going to have to do that the same way we develop trust with our human operators" and combat gear, Clark continued. "We will stress it to the point of breaking, we will understand why it broke, we will go back and fix that. We will iterate on it.

"We will develop a trust and an understanding of limitations – and where that trust ought to end," he said, "so the human beings can continue to make the best decision possible."

ATTACKING ARTIFICIAL INTELLIGENCE: How To Trick The Enemy

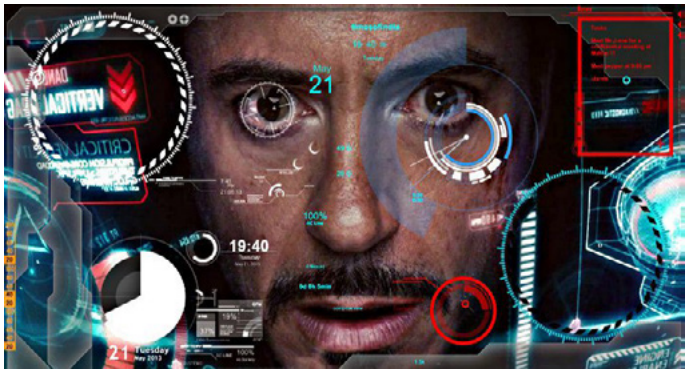

Iron-Man-Tony-Stark-jarvis

With the US, Russia, and China all investing in Artificial Intelligence for their armed forces, people often worry the Terminator is going to come to life and kill them. But given the glaring vulnerabilities of AI, maybe the Terminator ought to be afraid of us.

"People are saying, 'oh my god, autonomy's coming, Arnold is going to be here, he's going to be out on the battlefield on the other side,'" said Marine rifleman turned AI expert Mike Kramer. "I don't believe that. This is an attack surface."

As Kramer and other experts told the NDIA special operations conference (2/28/19-3/2/19) , every time an enemy fields an automated or autonomous system, it will have weak points we can attack – and we can attack them electronically, without ever having to fire a shot.

"If we're going to have an autonomy fight, have it at their house," continued Kramer, who now heads the technology & strategy branch of the Pentagon's Joint Improvised-Threat Defeat Organization (JIDO). "We are attacking the autonomy, not just the platform."

In other words, if you're worried about, say, the Russians' new robotic mini-tank, the much-hyped but underperforming Uran-9, don't dig in with your bazooka and wait until you can shoot at it. Use hacking, jamming, and deception to confound the algorithms that make it work.

How? Breaking Defense has written extensively about what we call artificial stupidity: the ways algorithms can misinterpret the world in ways no human ever would, because they interpret data in terms of mathematics and logic without instinct, intuition, or common sense. It turns out such artificial stupidity is something you can artificially induce. The most famous example is an experiment in which strategically applied reflective tape caused the AIs used in self-driving cars to misclassify a STOP sign as a speed limit.

But there are plenty of other avenues of attack, which is what Kramer & co. are talking about when they refer to "attack surface." At Carnegie Mellon University – home to the Army's newly created AI Task Force – a former Google VP turned dean of computer science, Andrew Moore, has come up with a simplified model called the AI stack, which shows how getting intelligent output from an AI depends on a whole series of underlying processes and technologies. Planning algorithms need models of how the world works, and those models are built by machine learning, which needs huge amounts of accurate data to hone its algorithms over millions of trials and errors, which in turn depends on having a lot of computing power.

Now, Moore devised the AI stack to help understand how to build a system up. But, his CMU colleague Shane Shaneman told the Special Ops conference this morning, you can also use it to understand how to tear that system down. Like a house of cards or a tower of jenga blocks, the AI stack collapses if you mess with any single layer.

The more complex and interconnected systems become, Shaneman continued, the more vulnerabilities they offer to attack. A modern Pratt & Whitney jet engine for a F-16 fighter has some 5,000 sensors, he said. "Every one of those can be a potential injection point" for false data or malicious code.

AI vs. AI

You can use your own artificial intelligence to figure out where the weak points are in the enemy's AI, Shaneman said: That's what DARPA's highly publicized Cyber Grand Challenge last year was all about. The stop sign tampering experiment, likewise, relied on some sophisticated AI analysis to figure out just where to put those simple strips of tape. This is a whole emerging field known as adversarial AI.

Machine learning uses arcane mathematical formulae called manifolds to extract patterns from masses of data. But no nation has a monopoly on math. If an adversary can see enough of the inputs your AI sucks in and the outputs it spits out, they can deduce what your algorithms must be doing in between. It turns into a battle between opposing teams of mathematicians, much like the codebreaking contests of World War II and the Cold War.

What's new, though, is it's also a battle of AI versus AI. One technique, called generative adversarial networks, basically locks two machine learning systems together in a virtual cage match, each driving the other to evolve more sophisticated algorithms over thousands of bouts. It's similar to the reinforcement learning system used in DeepMind's AlphaGo Zero, which played millions of games against itself for 40 days until it could defeat the greatest go players, human or machine. But generative adversarial networks add another layer. The two opposing AIs aren't identical, but diametrically opposite – one constantly generates fake data, the other tries to detect the counterfeits. What ensues is a kind of Darwinian contest, a survival of the fittest in which dueling AIs replicate millions of years of evolution on fast-forward.

One lesson from all this research, Shaneman said, is you don't want your AI to stand still, because then the other side can figure out its weaknesses and optimize against them. What you need, he said, is "algorithmic agility… constantly being adjust those weights and coefficients."

The good news is that the required combination of creativity, adaptation, and improvisation is a part of American culture – scientific, entrepreneurial, and even military – that potential adversaries will have a harder time copying than any specific algorithm. As former deputy secretary of defense Bob Workargued, Russia and China tend to see automation as a way of imposing central, top-down control and bypassing fallible human subordinates: The US military is looking at AI as a tool to empower human beings all the way down to individual pilots in the cockpit and junior non-commissioned officers (NCOs) in the trenches.

As rival militaries adopt AI, "they're going to accept more risk than the US is going to accept, and I think that at least initially… that's going to give them an advantage," said Nick Wager, an expert at the Defense Threat Reduction Agency. "But I think where the adversary will struggle is in the place he struggled in the past. It's the empowerment down at the NCO level, decision-making at the lowest level."

Russian Uran-9 armed unmanned ground vehicle

"Autonomy may look like an Achilles' heel, and in a lot of ways it is" – but for both sides, Wager said. "I think that's as much opportunity as that is vulnerability. We are good at this… and we can be better than the threat."

So don't fear the robotic reaper, Wager argued: "We can defeat that machine, which is, after all, easier to defeat than a human."

When It Comes to Artificial Intelligence, Ten Heads are Better Than One

By Todd Probert

The average consumer already uses AI on an almost daily basis, probably without realizing it. Mapping apps use AIs to direct people around accidents. Amazon uses AI to recommend your next purchase. AIs can even quickly read CT scans to diagnose cancer when human radiologists aren't available.

But we have reached a point where the AI capabilities being developed by commercial software companies can be applied to the battlefield, allowing commanders to make better warfighting decisions, particularly as conflicts occur across multiple domains.

And there are no shortage of companies interested in selling their AI algorithms to the military and intelligence community. However, nearly all of these companies are taking the same approach to developing AIs – going it alone without leveraging the knowledge and breakthroughs from outside their own walls.

The AI field is advancing and evolving so quickly, no single company can effectively tackle such complex technologies alone. The company that holds the magic key today may be leapfrogged by an upstart tomorrow.

Raytheon recognizes this, which is why we are leveraging the power of the crowd. We're building a broad and agile team of commercial tech leaders and innovators from outside the company — small and large firms from Silicon Valley, Cambridge and other centers of innovation — who are working side-by-side with Raytheon's own world-class talent to crack the greatest technological breakthrough since microprocessors.

The composition of this team constantly evolves as we identify new technologies and capabilities. From managing global logistics chains, to exploiting imagery data to identify potential threats, to assessing what adversary militaries may do next on the battlefield, our team is providing innovations from the best and brightest minds in tech today.

So where exactly did our team come from? In 2015, Raytheon began finding, partnering with and investing in innovative commercial companies developing AI, machine learning, analytics, autonomy and cybersecurity capabilities. Much like other government-led ventures you may already know, such as In-Q-Tel and DIUx, we're making investments in companies to accelerate their work, and then integrating those breakthrough technologies with our own to offer our customers truly advanced solutions.

Just as the military integrates the best technologies from a broad range of companies when they buy tanks, planes and ships, so too should they follow this model when seeking the best AI systems. Diversity drives innovation, and Raytheon's crowd development strategy fosters a diverse and growing AI team, ensuring the latest technologies are present in our solutions.

We know everyone benefits when we work together. Because when it comes to artificial intelligence, no one company can do it alone.

Todd Probert is Vice President for C2, Space and Intelligence at Raytheon Intelligence, Information and Services.

Raytheon

SIMULATING A SUPER BRAIN:

Artificial Intelligence in Wargames

Theater commanders around the world want weapons they can see and use right now, the Chairman of the Joint Chiefs told the Army War College. It's a lot harder, Gen. Joseph Dunford said, to sell experienced senior officers on an untested and intangible capability like Artificial Intelligence.

One week after Dunford's visit (April 2019), the Army War College convened two dozen officers and civilian experts to take on that challenge: How do you demonstrate the potential value of a military AI before you actually build it? (The conference was scheduled long before Dunford's visit, but his words were very much on participants' minds). The immediate objective: come up with ways to mimic the effects of an AI so the school's in-house game designers could turn it into either a computer simulation or a table-top exercise within 10 months — without new money. The hope is that the 2020 game, in turn, will intrigue Army leadership enough that they'll support a larger, longer-term AI effort.

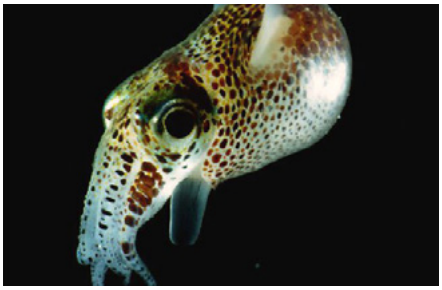

This squid's thought process is less alien to you than an artificial intelligence would be.

"We're trying to influence very senior leaders who don't have a lot of time," one participant said. "They just have to see it. Once they see it, the money will follow." (Sydney followed the Chatham House Rule in covering the event so no sources are identified by name.) That means, he said, those leaders need a way to visualize how AI might impact future operations. Wargames are a time-honored way for military professionals to see how new technologies might play out before you actually build them, often before it's even possible to build them.

The most famous case is the Naval War College in the 1920s and 1930s, which ran more than 100 games exploring a possible war with Japan, often with officers moving miniature ships around a tiled floor used as a giant gameboard. Admiral Nimitz famously said that these games explored so many different technologies and strategies that "nothing that happened during the [actual] war was a surprise… except the kamikaze tactics."

But wargaming Artificial Intelligence is a much harder problem. If you're trying to figure out how a new weapon might be used, the way the Naval War College did with the aircraft carrier in the 1930s, you can add new pieces to your board and new options to your rules: How does the game change if, for example, you let some of your ships launch aircraft x miles to bomb targets with y percent chance of destroying them?

But AI is not a physical weapon. AI is a machine that thinks. It's not too hard to change your wargame's rules to simulate planes that can fly farther, ships that sail faster, tanksthat are tougher to kill or satellites that transmit messages faster. But how do you change your game to simulate one side getting smarter?

Harder still: How do human game designers and human players simulate the military decision-making of a non-human artificial intelligence, when AI's crucial advantage is it can think of strategies no human ever could?

Game Changers

The Army War College doesn't have the time or money to build a superhumanly intelligent AI to play its wargames. Making such a mega-brain is probably years away for anyone. Even high-priority Pentagon AI programs are focusing, for now, on improving technical functions like maintenance, logistics, cybersecurity, electronic warfare, and missile defense — not building robot strategists.

You can't eliminate this problem, but there are two ways to reduce it. One is to simulate how future super-smart AIs will handle the staggering complexities of the real world by testing how today's relatively limited AIs handle the limited complexities of a simplified model. The other is to give human players some kind of advantage that helps

them out-think their opposition, simulating how warfare changes when one side consistently thinks faster and better.

It's remarkable how many ways the Army War College conference came up with to simulate AI without using any new technology at all. (I participated as an invited expert — with the Army covering my expenses — so some of these ideas are partially mine).

Even if your "simulation" is something as simple as two people playing a board game, you can break out a stopwatch or a chess clock to time their moves — and give one side more time each turn to simulate how AI can think faster than a human. You give one side more players, to simulate how AI can think through more options simultaneously and come up with a wider range of different strategies.

What if both sides aren't seeing the same game board? That's a common feature both of consumer strategy games, in which each player only sees the parts of the map around their units, and of formal War College exercises, in which each side is in a different room and gets all its information from a neutral umpire. This allows you to simulate AI by giving one side more or better information.

For example, one high-priority application for AI is rapidly collecting, analyzing, and disseminating vast amounts of data — say, video from surveillance drones — that would take human intelligence analysts much longer to plow through. You could simulate that capability by giving one side information faster than the other. Perhaps the umpire tells the humans playing the "AI-enabled" side where the pieces are on the board right now, but it only tells the non-AI team where the pieces were last turn, or two turns ago.

You can use a similar kind of lag to simulate how AI might automate a lot of routine staff work. It's an arduous process to turn a commander's scheme of maneuver into detailed timetables of what unit takes what route when, what air and artillery support they'll have on call, when they'll meet up with supply convoys carrying ammunition, food, or fuel, et cetera ad nauseam. It's also the kind of work which human's hunter-gatherer brains aren't evolved to do, but which computers are great at. So maybe the side simulating the AI-enabled headquarters can move its pieces on the board right now, but when the non-AI side issues orders, it takes a turn or two before its pieces actually carry them out.

Or you could combine both approaches. The side simulating AI gets to see the board and move its pieces then and there, but the non-AI side only knows where its pieces were last turn and can only issue orders for what they'll do next turn.

The Army also envisions AI being able to transmit intelligence and orders rapidly over secure wide-area networks, letting dispersed command posts work closely together despite the distance. You could simulate this by putting all the players on the AI-enabled team in the same room, while splitting the other team up. The division commander and his staff are in one room, his brigade commanders are each in their own rooms, and any communication has to go through the umpire. Now combine this with lag. All the AI players can communicate freely face-to-face, but if the non-AI commander tells his brigades to do something, they won't actually receive that order until the next turn.

Lockheed Martin's third Multi-Domain Command & Control (MDC@) wargame

If you can actually write or modify a computer wargame, instead of doing a pure tabletop exercise, your options get more sophisticated and fiendish. You can put each player at a different screen and manipulate how much information they get, how quickly, and how accurate it is. You can cut off in-game messaging between players to simulate radio jamming, or send false messages to simulate their network getting hacked. You can use crowdsourcing to simulate an AI's ability to generate a wider range of strategies. The non-AI players can only brainstorm among themselves, but the team simulating the AI can post the game board on the Internet and get hundreds of suggestions for its next move.

Once you put your wargame on a computer, you can also start replacing human players with real AI. Sure, it won't be as sophisticated as the future AI strategist you're emulating, but the simulated world it's playing in isn't as complicated as the real world, either. And there are a lot of AIs available already you can repurpose to play your particular game. "There's no need to reinvent game-solving AIs," said one participant. "That's already been done — and productized."

Those AIs are also getting better all the time. If you write your software right — specifically, the Application Programming Interface (API) — then you can plug in new and smarter AIs as they become available. Even in the long run, when you finally develop an AI strategist that could plan an actual battle, you'll really want to test it in simulation before you stake human lives on its performance.

RUSH TO MILITARY AI Raises Cyber Threats

As the US and other countries scramble to develop artificial intelligence (AI) solutions for military applications, the failure fix cyber vulnerabilities is teeing up a rush to failure, senior US and UAE AI gurus worry.

Frederick Chang, former director of research at the National Security Agency under President George W. Bush, told an Atlantic Council conference earlier this week that there just has "not been a lot of work at the intersection of AI and cyber." Governments are just "beginning to understand some of the vulnerability of these systems," he said. So, as militaries rapidly push to deploy systems they risk "increase the size of the attack surface" and create more problems than they solve.

Failure by governments to take proactive measures to ensure the security of AI systems "is going to come back to bite us," Omar Al Olama, minister of state for artificial intelligence for the United Arab Emirates, warned. "Ignorance in government leadership" is leading to deployment of AI "for AI's sake" — not because it is needed or is a wise thing to do. "Sometimes AI can be stupid," he said. Olama stressed that following the traditional commercial model of patching cybersecurity vulnerabilities after the fact would not work when building AI systems, because it "might be too late" for the security of nations and their citizens.

Chang explained that there are three major ways to attack machine-learning systems that researchers have not yet figured out how to thwart:

1. "Adversarial inputs" that can systematically fool a system's detector – something known as a "STOP sign attack" after an experiment in which researchers fooled a self-driving car by using masking tape to alter stop signs.

2. "Data poisoning" where an adversary might "alter data on which a system is trained" and cause its basic algorithm to reach wrong conclusions;

3. "Model stealing attacks" where adversaries infiltrate a system to figure out how to use its own operating system to thwart its functionality.

Col. Stoney Trent, chief of operations at DoD's Joint Artificial Intelligence Center (JAIC), agreed that education of leaders about the need to address cybersecurity in AI — and about the benefits and risks of AI in general – is needed. Another problem, Trent noted, is that there are few "testing tools and methods" to make sure AI systems work as they are supposed to and are not vulnerable to hacking. This is because in the commercial world spending time on testing is seen as a market risk, he explained. Thus, one of JAIC's tasks is to encourage development of such tools.

Cyberspace is one of the three "national mission initiatives" underway at JAIC, which stood up in June 2018 "to accelerate delivery of principally human-centered AI" across military mission areas. Trent said the effort "is not a place for the weak of heart," noting a number of barriers to his mandate to "accelerate delivery of human-centric AI" systems. These include technical barriers such as the need to "curate and categorize" data and proper problem scoping. The most difficult ones are not technical, but cultural. For example, he said DoD and service policies/practices regarding data sharing are a big problem. Another barrier is the tendency for development to take place in stovepipes resulting in bureaucratic resistance to cross-integration. "I haven't seen any evidence of it [integration] being done well in the military," he said wryly.