FIGHTING FOR SECONDS: WARFARE AT THE SPEED OF ARTIFICIAL INTELLIGENCE

John Brennan and Adarsh Kulkarni | 11.02.23



As the pace of artificial intelligence's advancement increases, US defense planners put a premium on the overmatching capabilities AI technologies portend for command and control. Significant planning, programming, and budgeting has commenced, which should produce increased capacity and survivability of communications on future battlefields. Now too are forces designing and prototyping modern systems built to help commanders sense more completely—observing,

FOLLOW US

FACEBOOK

YOUTUBE

TWITTER

DISCLAIMER

The articles and other content which appear on the Modern War Institute website are unofficial expressions of opinion. The views expressed are those of the authors, and do not reflect the official position of the United States Military Academy, Department of the Army, or

DISA

orienting, deciding, and acting—and faster than ever before. They will pierce the fog of war.

As timeframes of armed conflicts condense, what are the technical implications? Wars that might have unfolded over years in the past may be decided in months or even weeks. Operations executed over weeks must be completed in days or hours. And commanders who might historically have had the luxury of time before making a decision will be forced to do so in seconds. How will the organization and running of each individual command post change? These are the major questions facing military leaders as they chart a path forward that incorporates—and leverages the advantages of—autonomy, machine learning, trusted communications, and edge computing.

A Command Post on the Future Battlefield

2058 Hours, Local

A rabbit, chewing on some clover, freezes midbite and turns its ears to the left. Through the woods it hears the footfall of another creature moving toward it. Ready to flee if the source of the sound is a predator, the rabbit listens to the unidentified beast's four feet trot rhythmically and watches as it emerges at the edge of the clearing. It darts away from the approaching quadruped, which is unlike anything it has ever seen. The dog-like being is one of twelve in the area that night, autonomous robots that are taking up positions on key terrain surrounding the command post a motorized division is preparing to establish. These are the opening moments of a war and the dogs have been tasked with conducting reconnaissance.

Many minutes later, a low and steady clattering of armor treads emerges from the woods as the division commander and staff halt and establish a defensive posture. The command post's autonomous defense element—the dogs—now occupy security positions. The

Department of Defense.

The Modern War Institute does not screen articles to fit a particular editorial agenda, nor endorse or advocate material that is published. Rather, the Modern War Institute provides a forum for professionals to share opinions and cultivate ideas. Comments will be moderated before posting to ensure logical, professional, and courteous application to article content.

MOST POPULAR POSTS

Preparing to Win the First Fight of the Next War

The Return of Tactical Antiaircraft Artillery: Optimizing the Army Inventory for the Era of

dogs operate in groups of three. While one is on a four-hour shift, two others are recharging their batteries.

The dogs are proximate enough to maintain themselves as nodes in mesh networks they use for communication and data exfiltration. At this stage of the conflict, the dogs have one simple rule of engagement: alert the command post if there is any human approaching. If they detect a human, they engage their sensors to record and transmit audio, video, and thermal infrared data to the command post. A company commander assigned to defend the command post will decide what actions to take depending on whether the person is identified as friendly, enemy or noncombatant.

For this phase of the operation, the command post will remain in position for up to forty-eight hours before moving again. In those 172,800 seconds, the dogs are going to transmit a near-steady 288 kilobytes per hour, per dog. This is their command-and-control data rate, and it allows them to repeat a familiar pattern over and over again—look left, look forward, look right, report system status, repeat. This ensures they maintain communications with one another and keeps them alert. If they suspect or detect a human near them, then their data rate is going to swell to 720 gigabytes per hour of situational awareness observations from onboard sensors.

Bandwidth is expensive in this environment, in terms of power, weight, and cooling. There are also other claimants on the communications resources. Since every member of the division and its equipment is a sensor, there are additional intelligence, surveillance, and reconnaissance feeds competing for priority on the network. Logistical systems, voice communications, geospatial updates, command orders, medical diagnostic data, and weather data are all trying to move through these networks. Moreover, there are planned network restrictions and transitions from primary to alternate pathways in order to confuse any adversarial systems as to the true composition and identity of the division.

Small Drone Proliferation

Gaza's Underground: Hamas's Entire Politico-Military

UPCOMING EVENTS

There are no upcoming events.

ANNOUNCEMENTS

Announcing the Modern War Institute...

Essay Contest Call for Submissions: Solving...

Call for Applications: MWI's 2023–24...

Join Us This Friday for a Livestream with...

For these reasons it is important for the computer vision model in the dogs (or other sensors in the division) to have as few false positives as possible. Thankfully, the computer vision models aboard the dogs "eyes" have completed almost fifty thousand hours of training on highly diverse biomes, across all seasons of the year, and with countlessly diverse weather and light conditions. They also consumed special training data on the effects of modern military obscurants to vision, camouflage techniques and patterns, and differentiated gait and body poses of human and robotic military formations. As impressive as this training data regime sounds, it can never account for all circumstances, which is why the division's machine learning operations element is prepared to fine-tune the dogs and other sensors across the force on the biomes they encounter and the inevitable change in enemy tactics, techniques, and procedures that will emerge from the first battles.

Retraining a computer vision model is like all other statistical analyses. Simply speaking, there needs to be a representative sample to express mathematically the relationships among the sample's key variables. In computer vision, the objects in a scene—the sky, clouds, the horizon, hills with trees, hills with no trees, individual buildings, windows on buildings, faces looking through the windows, and so on—are the sample. Each dog's sensors are collecting countless samples of the biome (the population) surrounding it. Within these samples are new outliers the computer vision model has never seen before, because its makers did not have a laboratory that could contain all the diversity on Earth.

After the first 8,640 seconds of night patrol, the dogs have falsely alerted on a bear that lumbers up to a tree and then stands on its hind legs to try to reach some acorns. When it does this, one of the dogs detects it, captures several minutes of it via its sensors, and alerts the command post security element. At the morning staff huddle, the command post security company commander asks the machine learning operations element what they can do about the bear that keeps them awake at night.

The machine learning team is one of many across the joint force. They are all gathering samples in the early stages of the conflict. The perceived outliers, as well as numerous random samples of all observations, are flowing back to compute clusters in and out of the theater based on their capacity and queue priorities. Some of the data travels over the tactical mesh networks. Other data is passively gathered by the various logistics vehicles moving from unit to unit, and then moved via physical storage device transfers.

As data scientists work through the false positives across the entire visual sensor fleet, they determine they have statistically significant observations of the local bear species that they can pass all the observations back to AI data product teams to complete a round of semisupervised negative training to classify them as nonhuman and noncombatants. The teams believe the newly trained computer vision model will have around 70 percent fewer false positives for bears.

Another important finding in the early samples is the discovery of an enemy tactical sentry that is not being detected by the computer vision model. This false negative could have a detrimental impact on force protection. The sentry resembles a bird sitting on a tree limb, but really it is a discrete audio and visual sensor package. By comparing multiple observations of the sentinel to a sample of videos from amateur bird watchers on the internet, the data science team is able to train a classifier to distinguish the sentinel from live birds.

Armed with this new computer vision model, the team prepares its software update strategy. The computer vision update is one of hundreds that will be deployed in the next 21,600 seconds across the joint force. They cannot nor should they all be deployed at once. The teams have practiced this continuous integration and continuous delivery (CI/CD) approach hundreds of times in the last three years. Like modern cloud computing providers, they have a deployment strategy that goes

through a series of tests and small deployments to confirm proper functionality.

In addition to these modern software practices, the machine learning operations elements are also ensuring their computer vision model deploys to a sample of dogs that will reobserve the bird-like subjects and where the bears generated the most false positives. With fresh data showing the false positives and false negatives are significantly reduced, the teams continue their deployments across the rest of the fleet of computer vision sensors. Thankfully, updating a computer vision model is efficient on bandwidth. The original model is forty to sixty megabytes depending on the sensor. Updates can be 10–20 percent of the original model, which in this case is just five to ten megabytes. All these little efficiencies matter across the tactical edge and its potential for denied, disconnected, intermittent, or limited conditions.

Data Epochs, Battle Drills, and WERX

This vignette describes one of hundreds of updates that could happen every two weeks—1.2 million seconds—in the onset of a conflict. The series of updates that must occur during each of those data epochs are the new battle drills. They must become as familiar and honed as a soldier clearing a jammed weapon.

What concepts must be demonstrated for eventual fielding? How must our military formations, personnel, and leaders adopt and mold these new data battle drills? The US Army has started a digital odyssey into cloud computing, software-defined networks, algorithmic warfare, autonomy, and human-machine teams. Each of the US military services is rapidly enabling its acquisition, security, and IT cadres to deliver these technologies and practices. The resulting software development practices are not limited to back-office functions. They are emerging simultaneously at all echelons through traditional and alternative acquisition pathways. They are being accelerated by tried-and-true research,

development, test, and evaluation programs, as well as new organizations singularly focused on successful fielding of digital and artificial intelligence capability. Ushering in these new efforts are a network of leaders who share a common vision and a renewed appetite for prudent risk-taking.

Walk into any of the new software factories or Dev/Ops (development/operations) software teams—like AFWERX, SOFWERX, and others—and you see empowered professionals who understand their mission and craft. Like the pioneers before them in aviation, mechanized operations, and uncrewed systems, they know the novelty of what they are undertaking can easily be misunderstood or dismissed. Yet as they contemplate the ongoing wars in Ukraine and Israel, they energetically push for new capabilities to overmatch any opponent.

A Roadmap for AI Acceleration

All of these efforts are meaningful. For them to bear fruit and maximally impact military effectiveness, leaders must actively work to accelerate their teams' effects. There are several key steps they can take to do so.

First, remove the distance between development teams and their end users. This is crucial to successful software and hardware adoption. Regular touchpoints among the users and developers is something you should measure and invest in. If the users are not breaking the technology, and the developers are not rapidly adapting it to the use cases they could not envision, then you are deferring those unfortunate discoveries to your first battle.

Second, fight the status quo bias in the bureaucracy. If anyone tells you that there are no requirements for generative AI, remind them that there also were no requirements for firearms, tanks, planes, counter–improvised explosive devices, and counter-drone technologies at the other inflection points of military

DISA

history—until technology and use cases converged and made the requirement abundantly clear.

Third, break down your AI program into a portfolio that balances a general development, security, machine learning, operations (DevSecMLOps) pipeline with the unique needs of different machine learning disciplines (e.g., natural language, computer audition, and computer vision) for different applications like remote sensing, sense-making, and autonomy.

Fourth, in addition to the human personas your teams are building for, you must bring another user to the table—the machine personas in your emerging force structure. Each machine learning–enabled device in your force structure will need an AI-ready data product strategy and a data networking strategy to enable its retraining across a spectrum of conflicts, biomes, and missions. While the humans continue to consume beans and bullets, the machines will consume batteries and bytes.

Fifth, beyond the physical and virtual training environments you have already built, you need a device farm, a model zoo, and a family of digital twins. Mobile application developers have to test across a variety of device manufacturers to ensure forward and backward compatibility. Device farms help developers quickly get technical feedback from such heterogeneous device fleets. In the next era of warfare, these might be quadrupedal robots, drone swarms, or optionally crewed tracked vehicles. Machine learning engineers build and iterate on numerous open source and proprietary models. As they train and retrain them, the engineers will have a growing collection (zoo) of models that serve as a better jumping off point for quick fine-tuning in a novel situation. Digital twins will help ensure you have an adequately representative environment for showing your new network-device-data-model combinations work as intended.

Sixth, trust the Zero Trust movement. Like the frequency-hopping spread spectrum invented by Hedy Lamarr and

George Antheil, zero-trust systems constantly enforce credential and other certificate connections, rotations, and reauthentications to ensure security consistency against outsider and insider threats. Getting it right is a complex art form, but it is a capability you should demand across your data and system accesses. BYOD (bring your own device) is an indicator your organization can operate in a zero-trust manner. When you can trust the heterogeneous hostile foreign devices brought to work by your employees, you will achieve the organization and technical prowess needed to do the same with your allies and mission partners in the combined, joint force.

Seventh, your organization needs to deploy into production every week and your leaders should participate in the CI/CD process. Modern cloud computing services achieve high availability through several strategies. Key among them is the software-defined connection between the developers and their leaders during an outage event. In a pure form, DevOps means the same team that built the software also deploys it and will be paged in the middle of the night if the deployment fails. This enhances attention to detail during testing. The paging process can also be connected to key leaders. If you know you have a certain amount of time to fix the bad software deployment and return to normal services before your manager gets paged, then everyone pays more attention to detail during testing.

Eighth, clarify the areas where you will allow experimentation on “System 1” and “System 2” machine thinking among your various capabilities. In his book *Thinking, Fast and Slow*, Daniel Kahneman popularized the modern understanding of the thinking abilities of humans as being the by-product of two systems. System 1 responds automatically and quickly—like a fight-or-flight response. It takes little to no effort for your senses of sight, hearing, or smell to process a house on fire. Kahneman distinguishes this from System 2, which “allocates attention to the effortful mental activities” and enables tasks like complex computations. Through it, we

feel much more agency and choice over when we allocate our energies to concentration. The key to human-machine teaming will be the allocation of System 1 and System 2 tasks to the right teammate.

Forbid the day we experience this, but on a future battlefield two armies will clash in the night. In the morning, several battalions on each side may be mere remnants of themselves. Whichever side can execute a reconstitution battle drill faster will have the immediate advantage. The mission order to detach from a home organization and attach to a new one requires several administrative and control steps. The commander should choose which new leader and which formations to assemble first—this is a System 2 process. Once the order is issued, though, a series of System 1 applications and messages should enable all the personnel, medical, logistics, and other supporting data to swiftly and accurately transmit to the new command staff. If we do the work that is before us well, that new battalion should be ready to fight in 3,600 seconds rather than 28,800. Every second counts.

John Brennan, PhD is the general manager of Scale AI's public sector business. He served in the US Army as an infantry officer in Korea. He also served in a war zone with the CIA. Over the past twenty-five years, he has helped governments worldwide develop and operationalize these technologies for their cloud computing, data science, and cyber security programs. John was the product manager for AWS's secret and top secret clouds supporting the intelligence community and DoD. The article is inspired by his work with Army Futures Command and its efforts to design future CJADC2.

Adarsh Kulkarni is a senior software engineer at Scale AI, where his work focuses on computer vision for autonomy. With degrees in mechanical engineering and robotics, he previously worked as an autonomy and AI applications engineer for Ghost Robotics. The article leverages Adarsh's

knowledge in operating robotic dogs and his work at Scale AI helping train computer vision models for perimeter defense applications.

The views expressed are those of the authors and do not reflect the official position of the United States Military Academy, Department of the Army, or Department of Defense.

Image credit: Sgt. Darrell Stembridge, US Army

LEAVE A REPLY

Your email address will not be published. Required fields are marked *

COMMENT

Save my name, email, and website in this browser for the next time I comment.

NAME * EMAIL * WEBSITE

POST COMMENT

© 2024 Modern War Institute