



SERIES
Number 6

INTELLIGENCE AFTER NEXT

ENSURING DECISION ADVANTAGE ON THE FUTURE BATTLEFIELD—
INTELLIGENCE AT THE SPEED OF HYPERSONIC WARFARE

by Joseph P. Convery

Intelligence in the Age of Hypersonics

Hypersonic weapons are adding a new dimension to the pace of warfare and will drive engagements on the battlefield at blistering speeds. This will require military commanders to act faster than the advanced weapons and automated processes available to their adversaries.

Achieving decision advantage in this operational environment must begin with the intelligence activities that underpin all military operations. Timely and accurate intelligence offers the information advantage that enables the decision cycle. Applying automation to aspects of the intelligence cycle, and building trust in those processes, will enable the sensor-to-shooter constructs essential to interdict advanced weapons and meet increasing operational demands for timeliness.

Failing to meet the demand for timely intelligence will result in the loss of decision advantage on the battlefield, the subsequent loss of the operational initiative in combat, and potentially the fight.

Artificial intelligence (AI)-based solutions will offer a variety of advantages on the battlefield and across the intelligence cycle. When combined with resilient Intelligence, Surveillance and Reconnaissance (ISR), and advanced analytics, it will deliver unprecedented capabilities to operational forces. However, AI alone will not fully address this challenge. We must build trust for the intelligence derived from AI processes, for both the consumer and across the Intelligence Community(IC). Trust is the key to enabling their use, and as such our ability to reap the full benefit of automation.

Introduction

The speed of combat continues to increase, with hypersonic missiles capable of striking targets hundreds to thousands of miles away within minutes of launch. Their speed, flight profiles, and maneuverability complicate detection, and will limit both warning and

reaction time, and thus a strategic or operational commander's decision window. However, these weapons are just the more visible examples of developments that will speed the pace of engagements during conflict. Automation applied by enemy forces to their command and control, targeting, deception strategy, and non-kinetic capabilities will also severely restrict U.S. commanders' ability to grasp the intentions of their enemy, and to act first to seize the operational initiative. The challenge to secure decision advantage will rely largely on the IC's ability to predict an adversary's likely course of action and provide the earliest detection of those enemy activities that could rapidly alter the course of battle.

Intelligence, by its nature, is a cycle that includes planning to address a specific requirement; the collection of raw information from a variety of sources to address that need; the exploitation of that information to synthesize valuable data relevant to that requirement; the analysis, integration, and evaluation of that information to craft a finished product; and the dissemination of that product to the appropriate decision maker. The cycle begins again as the intelligence product is consumed and new or additional requirements are articulated by the consumer. For the intelligence consumer at the strategic, operational, and tactical levels of warfare, the quality of information and speed at which that intelligence is provided shapes, and will continue to shape, the course of battle.

The application of AI to specific aspects of the intelligence cycle may allow the IC to fulfill its critical responsibilities within the decreasing time available for commanders. However, this change will not rest solely on the application of technology but will also require significant changes within the human factors that drive the commander's decision cycle, as well as those that drive the collection and analysis of the intelligence critical to those decisions. None of this will be easy.

Artificial Intelligence as Decisionmaker

As advanced capabilities continue to shorten the decision cycle, the speed at which knowledge is developed and intelligence is then delivered to the decision maker must increase, while maintaining the quality of that information. This reality has not been lost on the IC, having already sought to improve the speed of collection, processing, exploitation, and analysis via the application of artificial intelligence and machine learning (AI/ML) processes. There is now recognition that the speed of “human in the loop” processing and dissemination must substantially increase to preserve the time-sensitive value of perishable data, such as the location of mobile missiles and other battlefield transients.

Going forward AI will drive significant change in the character of warfare, continually increasing the pace of operations on the battlefield. AI applications will enable both friendly and adversary leadership to direct operations and weapons employment faster, as each side seeks to control the operational initiative on the battlefield, a key element of achieving victory.

But what happens when the pace of AI driven combat exceeds the commander’s ability to control the battle? With both sides employing AI to drive weapons applications the choice to slow the pace of conflict won’t be solely that of the friendly commander, and neither side will want to cede the initiative to the other. Each commander will be forced, to some degree, to trust in the decision making of their machine-based counterpart.

As a result, the speed of war is driving the inexorable need to extract some parts of human interaction from the sensor-to-shooter process or risk being “beaten to the punch” by the AI-driven processes and advanced high-speed weapons of an adversary.

While information must be rapidly disseminated from sensor to shooter, the decision to “put steel on target” without human review, or to make other force

applications without command deliberation, will be anathema to most leaders. As such, trusting the decision of a machine during a peer level conflict, to drive what will likely be multiple, near real-time targeting decisions at critical points in battle, may be difficult for commanders.

During the Army Futures Command test of AI-enabled targeting capabilities, known as Project Convergence 2020, the Army stated that AI and autonomous capabilities have decreased the sensor-to-shooter timeline from 20 minutes to 20 seconds.¹ Brigadier General Ross Coffman, Army Futures Command, stated that “the technology exists, to remove the human, but the United States Army, an ethical based organization, is not going to remove a human from the loop to make decisions of life or death on the battlefield, right? The artificial intelligence identified and geo-located enemy targets. A human then said, Yes, we want to shoot at that target.”²

While this was one engagement, in a test environment, it demonstrates the potential power of AI in a targeting application. This also illustrates a belief that all decisions of this magnitude embody an ethical requirement for human direction, and that even in the heat of battle, decisions of this type will demand human approval. However, if the technology exists to do so without human intervention, and the pace of operations demands action, what will be required for a battlefield commander to trust AI?

Can the intelligence underpinning an AI-driven decision to initiate a strike be the key to that commander’s trust... and what if that intelligence itself emerges from an AI-based effort?

In a similar sense, this internal conflict in accepting machine-driven analytic solutions will prove difficult for most intelligence professionals. Intelligence analysts are trained to question their assumptions, seek alternatives, and apply all aspects of analytic tradecraft to drive their

analysis toward increasing levels of confidence. Their actions to achieve higher levels of confidence are, in fact, designed to engender trust in those that would use their analysis as the basis for critical operational decisions.

So, at the advent of creating the automated solutions critical to enabling success on the modern battlefield, we have both the principle consumer of intelligence, and the community that generates it, inherently untrusting of solutions devoid of human context. As we seek to increase the speed of generating information within an increasingly constrained decision cycle, we must find a way to engender trust in that process or risk being rapidly overwhelmed by the AI-driven capabilities of our enemies.

Trusting a machine

As we develop automated technologies for the warfighter, we must also build their trust in the machine-driven intelligence solutions that will feed those capabilities.

U.S. Army doctrine describes trust as a value that is “gained or lost through everyday actions more than grand or occasional gestures. It comes from successful shared experiences and training, usually gained incidental to operations but also deliberately developed by the commander.”³ This description would imply that trust can be earned by an automated process based on the dependable and consistent performance of that procedure.

We trust the global positioning system (GPS), for example, because it performs reliably every day. Even though we don’t see and prove the thousands of calculations playing out on every trip, we trust it because it has proven to be reliable and consistent over time.

This trust is often strengthened by our faith in the skills of the experts who developed the capability, their understanding of the technology, and their ability to understand and address shortfalls in its performance during the development cycle.

ARTIFICIAL INTELLIGENCE MUST BECOME COMMONPLACE IN ALL ASPECTS OF THE INTELLIGENCE CYCLE AND BE RECOGNIZED AS A TRUSTED TOOL CAPABLE OF SUPPORTING THE FULL RANGE OF INTELLIGENCE FUNCTIONS.

In the past, nascent AI capabilities focused on performing general, repetitive tasks, achieving a lower error rate than their human counterparts. The value of AI was first recognized in this way, building an initial sense of trust for these rudimentary capabilities. AI was good in this application, and there was only limited risk to its use.

From a developer’s perspective, their trust in the AI they developed was further buttressed by the predictability and transparency of the algorithms and processes the system used to generate solutions. Its purpose was clearly defined and limited within the algorithm, as were the situations when it should be used.

In more risky applications where legal or ethical considerations were in play, developers sought to apply multiple, varied AI algorithms to automatically compare results and identify error rates beyond acceptable parameters for risk.

AI systems in use today are engineered to dependably provide accurate results for specific tasks, and while dependability and accuracy were the actual goal, the cornerstone for human trust was set by the quality of each development effort.

To enable the combatant commander to hand over even a portion of his or her decision authority to a machine will require proven performance of that automated process over time. It will require trust.

The same will be true for an intelligence analyst working with the results of an AI algorithm to assist in the creation of an analytic position. Without developing trust in AI processes, and the data underpinning those algorithms, we risk human-generated delays that could lead to mission failure.

However, it is no longer a question of whether a battlefield commander will have to trust machine driven solutions and decisions to act, but one of how to build that trust. Because AI-driven processes are a reality for the battlefield commander, they must also become an accepted reality for the IC or the IC risks becoming the “choke point” for the critical information that fuels a more pressing decision cycle.

Engineering Trust

‘Train as you fight’ is a military philosophy equally applicable to the use of AI. To effectively engineer trust into both intelligence activities and the combat operations that they enable, they must be subject to rigorous testing and progressively more complex application during peacetime, when analysts and operators alike will have time to compare human results with the findings of their new machine-driven partner. As those comparisons consistently demonstrate reliability, the general need to question those results will subside.

Securing AI from malicious interference will itself also be a key element of trust as we develop faith in the performance of AI processes. Regular sampling and data comparison must be accomplished to ensure the algorithms themselves, or the training data used in their development, have not been tampered with, or otherwise impacted for malicious purposes.

How do we engineer trust as part of the development cycle for an AI-based intelligence process? Engineers are adept at validating the performance of a developmental capability, repeatedly testing, improving, and retesting each aspect of that capability, under a full range of environmental and operational stressors.

This helps determine limitations and points of failure for that capability, leading to improvements and increased performance as required by the operational need.

For an AI system, its behavior, or what we see as its performance, can also be improved by the experience it gains performing a specific task. As the AI system applies its inherent mechanisms for adaptation and learning, it should grow more proficient and reliable over time.

Engineering “trust” in an AI-driven analytic solution can be done the same way, but with the added advantage of having the state of the art machine the process seeks to emulate, the human analyst, sitting right next to the developmental capability to serve as an example for its desired end state.

Involving the stakeholders, analysts, and decision makers alike in the development process will not only assist in the functionality of the AI process, but build trust in the very consumers who will rely on it.

In this case, it is also important to understand the rate of human error. For example, how often does an analyst make a simple typographical error such as transposing a digit in a latitude/longitude entry in a structured database? Can an AI process eliminate that error? Can AI be used to eliminate normal human rates of error, and what level or type of errors will be acceptable within an AI driven analytic process?

The direct interaction of the algorithm with human analysts has improved its performance, but those analysts now have seen the machine replicate their own abilities, at least for a limited set of training data. As they continue to review results based on new data, and find those results satisfactory, they will develop a greater trust in the application.

Once comfortable with this capability, they will realize the time savings it offers and accept the process, with that process becoming a matter of routine. The machine will have achieved, or rather earned, human trust.

Trust in AI can also be improved through the hierarchical satisfaction of requirements. This is to say that having a hierarchy of automation needs, where each increasing level of task is met and proved to be accurate and reliable, with successive evolutions building upon one another, can engender that trust in performance of each new capability. Low-level tasks that can be automated at high accuracy today can help with a trust roadmap, as we proceed to more and more complex AI processes.

The Intelligence Cycle: Seeing, Understanding, Sharing, and Acting Faster

Intelligence analysts are challenged to gain meaning from the wealth of raw information available to them in peacetime. As intelligence collection capabilities have improved, this challenge has become progressively more difficult in terms of volume, variety, velocity and verifying the veracity of that data. The shortened timelines and pressures of major combat operations, such as those likely in combat with a peer competitor, will significantly increase this challenge while reducing the time available for the IC to meet critical wartime demands. The IC must literally see, understand, and share intelligence faster to enable the operational consumer to act faster.

Automating aspects of the intelligence cycle must be accomplished to reduce the time from discovery to enabling an effective combat action. Automation can be applied to detect and understand a threat faster, share that information more rapidly, likely across multiple security domains, and enable the combat commander (or his AI surrogate) to act quickly to neutralize that threat and maintain the operational advantage on the battlefield.

The IC's support to each of the Combatant Command operational plans are underpinned by a complex "multi-INT" collection strategy, comprised of numerous collection requirements working together to meet that commander's needs.

High intensity peer-level conflict will draw most US collection capabilities to the fight and further challenge the human aspects of the intelligence cycle. While facets of automated collection queuing are already available, the scope of this challenge in wartime must be addressed to resolve what will certainly be an overwhelming wealth of information which must now be quickly digested to support the fight.

How we apply this technology to each portion of the intelligence cycle, weave those processes together within a coherent architecture, and allow intelligence gained from this process to be ingested into other AI-driven processes on the battlefield, will be key to the successful integration of AI.

Seeing Faster

To meet the challenge of speed within the collection portion of the intelligence cycle, we can exploit emerging technologies to benefit several activities, including collection management, exploitation, and dissemination of GEOINT.

The challenges associated with GEOINT provides a strong example of where AI can offer significant benefits within the collection portion of the cycle.

Robert Cardillo, former Director of the National Geospatial-Intelligence Agency, captured the imagery exploitation challenge perfectly in his statement regarding the expansion of commercial imagery sources: "If we were to attempt to manually exploit the commercial satellite imagery, we expect to have over the next 20 years, we would need eight million imagery analysts."

Within the context of GEOINT support in wartime, the ability to rapidly detect and identify specific targets by type will be a crucial supporting task for an increasingly rapid targeting cycle. If a mobile missile can be detected and identified in the brief window where it may remain stationary to prepare for and

conduct a launch, the launcher itself can be destroyed, preventing reload and re-use.

To achieve the capability necessary in this scenario will require an AI-enabled intelligence architecture that supports the rapid identification of the target, and an ability to pass required targeting information quickly to the “shooter.” On-board sensors, or rapid ground-based processing of imagery-derived data, capable of geolocating specific targets by type, and then injecting a required data set into the targeting process must be the goal of support to combat operations.

While the GEOINT collection example provides a strong case for the benefits of AI, the same or similar strategies will likely be applicable to collection activities within each of the intelligence disciplines. Collection management tasks, “Cross-INT” collection cueing, automated geographic and temporal correlation of data derived from different sources, and other collection and exploitation functions can be supported by AI to speed aspects of the collection and exploitation portions of the cycle.

Understanding Faster

While AI, and automation in general, can be applied to speed collection tasking and the detection of threats, improvements within the collection portion of the cycle may become a double-edged sword. The immediate benefits offered by applying AI to collection management and target recognition will support addressing immediate operational needs, however, improved collection will increase the amount of information that must be consumed by the analytic corps as they seek to derive meaning from an increasing data flow.

This is further compounded by the addition of “Big Data,” as analysts seek to derive meaning from massive data sets that may offer unique insight into a specific

area of interest. The challenge for the analyst becomes assimilating this data while maintaining accuracy and achieving the timeliness required by their operational consumers.

Artificial Intelligence and automation in general must be applied to assist the analyst in drawing value from the massive amounts of data that will likely be available to them during combat with a peer competitor. To consume the raw intelligence necessary for an analyst to provide an assessment or support the commanders targeting strategy, AI/ML capabilities are already offering some support, but how can AI be applied to ease the burden on the analyst, prioritize key pieces of information critical to their area of interest, and offer the time savings needed to dive more deeply into their key questions?

“Big Data” is one analytic challenge well-suited to the application of AI, to sift through, sort and present the truly relevant portions of that data to the analyst, and it can do so in a fraction of the time required for human interaction with that data.

AI can also reduce the data the analyst must actually look at through triage type processes (assigning priorities), duplicate detection, and data aggregation, leaving only the data most relevant to the task for the analyst to consume manually.

But what about the fog of war for both the battlefield commander and intelligence analyst alike? What happens when the speed of change and the scale of observables combine to be beyond our ability to effectively grasp? Can strategic reasoning supported by AI help us to maintain situational awareness when the pace and complexity of change, like that which will be faced in modern warfare, exceeds human comprehension?

Military leaders and intelligence analysts both rely on years of experience to form their perceptions of a potential threat. They make decisions or draw

conclusions based in part on their personal perceptions of the information they have at hand.

In high pressure situations, experience is critical for an effective leader to rapidly shape a decision. However, in a dynamically evolving situation, when massive amounts of information must be considered to arrive at the most effective conclusion, this can easily become overwhelming.

Humans use various approaches to manage and digest large amounts of data, but human judgment remains the primary factor for making decisions. If they fail to recognize information hidden within the massive amounts of data, or react to that data with a mistaken bias, their decisions could become unsuccessful in achieving intended goals.

AI may offer some relief. Without bias, AI may offer more effective decision alternatives, seeing trends within an enemy's actions perhaps unnoticed by the analyst or battlefield commander. The human, then, is not forced to comprehend a tidal wave of data and alternatives but interacts with AI-driven alternatives based on its unbiased interaction with the data.

As the IC has recognized the advantages of AI, the challenge next becomes building the architecture to allow analysts to take advantage of it. This must include the ability to share the data derived from AI with the right operational consumers and in the time needed for them to address the threats discovered within this progression.

Sharing Faster

Within the intelligence cycle, dissemination forms the key linkage to the consumer, giving the intelligence gained true operational value. Over time, various strategies were attempted to offer the decisionmaker the information they needed in the right form, and at the right time, to support their conclusions. A push strategy often overloaded the consumer with more data than they required or could assimilate into their decisions, forcing

the consumer to spend time extracting the data most important to them. A pull strategy sometimes created additional tasks for the decisionmaker to go look for the data they needed, further complicating their efforts.

Technology may offer time savings by enabling rapid dissemination of the most important information to a specific consumer based on the specific needs of that person, their mission set, their location, and other variables as needed to refine the reporting delivered to them.

Another key intelligence dissemination challenge for the IC has been pushing classified data between domains. A similar construct can be created to serve as a gatekeeper to prevent data at higher levels of classification from being disseminated on networks with more restricted classification levels. Intelligence used to locate a mobile missile may be held at a higher classification than the network supporting the strike aircraft can support. Requiring a human intensive process, either to vet each message or re-draft/tag reporting for additional dissemination, would take time and risk relocation of the target prior to the attack.

Acting Faster

AI-driven strategic reasoning, as discussed earlier, can be applied to massive data sets in the heat of battle to present more accurate, unbiased alternatives to our combat commanders and intelligence analysts alike. This capability will buttress an already skilled operator, enabling clarity of situational awareness while providing a range of potential alternative solutions, further increasing the effectiveness of their subsequent decisions.

AI-machine learning approaches, such as Reinforcement Learning (RL), can be used to create algorithms that will recognize trends hidden from human view, and do so rapidly across massive amounts of data changing moment to moment as the pace of battle increases. Offering alternatives, without the

**IT IS BEST THAT WE SEEK TO MASTER
ARTIFICIAL INTELLIGENCE NOW OR BE
MASTERED BY IT TOMORROW IN THE
HANDS OF PEER COMPETITORS ON THE
BATTLEFIELD.**

ingrained biases of an individual's personal experience, will provide more effective paths to achieving the individual's goals, be it a combat objective or understanding the nature of an emerging threat.

As AI-driven intelligence findings earn the trust of the operational consumer, we must continue to seek their direct engagement to enable the emerging AI processes to drive friendly force applications.

Conclusion

The employment of AI is no longer a question of if, but of how fast we can begin to take advantage of the improvements it can offer across the intelligence cycle. To do this we must address both our intelligence professionals and our consumers' need to develop trust in the AI processes that will underpin modern warfare.

We must then design the overall architecture that will allow AI enhancements to be rapidly incorporated into existing processes, interface with other AI capabilities as they come online, and share data at the fastest possible speed.

It is important to note that AI development today in more complex applications is simply difficult to do well. Object recognition provides a strong case study for the difficulties associated with developing an AI capability to replicate a function that requires a great deal of uniquely human expertise and skill.

As we move into an AI-driven battlespace of the future, it will take time, commitment, and resources to create AI capabilities that will meet those standards which make them worthy of trust. However, as AI development continues to improve, the ability of an AI application to learn and evolve may provide its own path toward achieving those standards. It will learn and improve, perhaps earning greater trust.

If the IC is to remain relevant to the fight, we must be capable of delivering information advantage in the brief window the commander may have to decide the course of battle. The commander must then also be prepared to act on that information at the speed required to disrupt an adversary's own decision cycle. AI is currently the best hope for achieving the speeds necessary to meet these requirements.

While human intellect will still be critical to the strategy of war, we must accept that some decisions must fall to the responsibility of AI or risk defeat by an enemy's more rapid application of force.

AI's nascent employment in unique applications is only a glimpse of its future utility; it ultimately will offer advantages across every element of the intelligence cycle. As such, continuing to limit AI for such unique, standalone applications, only used by seasoned data scientists, will continue to delay its potential for positive impact. AI must become commonplace in all aspects of the intelligence cycle and be recognized as a trusted tool capable of seamlessly supporting the full range of intelligence functions.

Because the capability of AI is growing exponentially across the spectrum of warfare, it is best we seek to master this capability now or be mastered by it tomorrow in the hands of our peer competitors on the battlefield.

Endnotes

1. C4ISRNet, Inside the Army's Futuristic Test of its Battlefield Artificial Intelligence in the Desert, Nathan Strout <https://www.c4isrnet.com/artificial-intelligence/2020/09/25/the-army-just-conducted-a-massive-test-of-its-battlefield-artificial-intelligence-in-the-desert/>
2. Ibid.
3. Army Doctrine Reference Publication (ADRP) 6-0, Mission Command, Washington, DC: U.S. Government Printing Office, May 2012, para. 2-5

Author

Joseph Convery is a Principal Intelligence Collection and Technology Engineer. He is a former U.S. Army Intelligence Officer, and civilian collector and analyst with over 37 years of experience across the defense intelligence community. He currently serves as a special advisor on Combatant Command intelligence requirements and programs, and in support of missile defense left through right of launch integration activities.

Intelligence After Next

MITRE strives to stimulate thought, dialogue, and action for national security leaders developing the plans, policy, and programs to guide the nation. This series of original papers is focused on the issues, policies, capabilities, and concerns of the Intelligence Community's analytical workforce as it prepares for the future. Our intent is to share our unique insights and perspectives surrounding a significant national security concern, a persistent or emerging threat, or to detail the integrated solutions and enabling technologies needed to ensure the success of the IC's analytical community in the post-COVID-19 world.

About MITRE

MITRE's mission-driven teams are dedicated to solving problems for a safer world. Through our public-private partnerships and federally funded R&D centers, we work across government and in partnership with industry to tackle challenges to the safety, stability, and well-being of our nation.