Operationalizing Artificial Intelligence for Algorithmic Warfare

Courtney Crosby, PhD

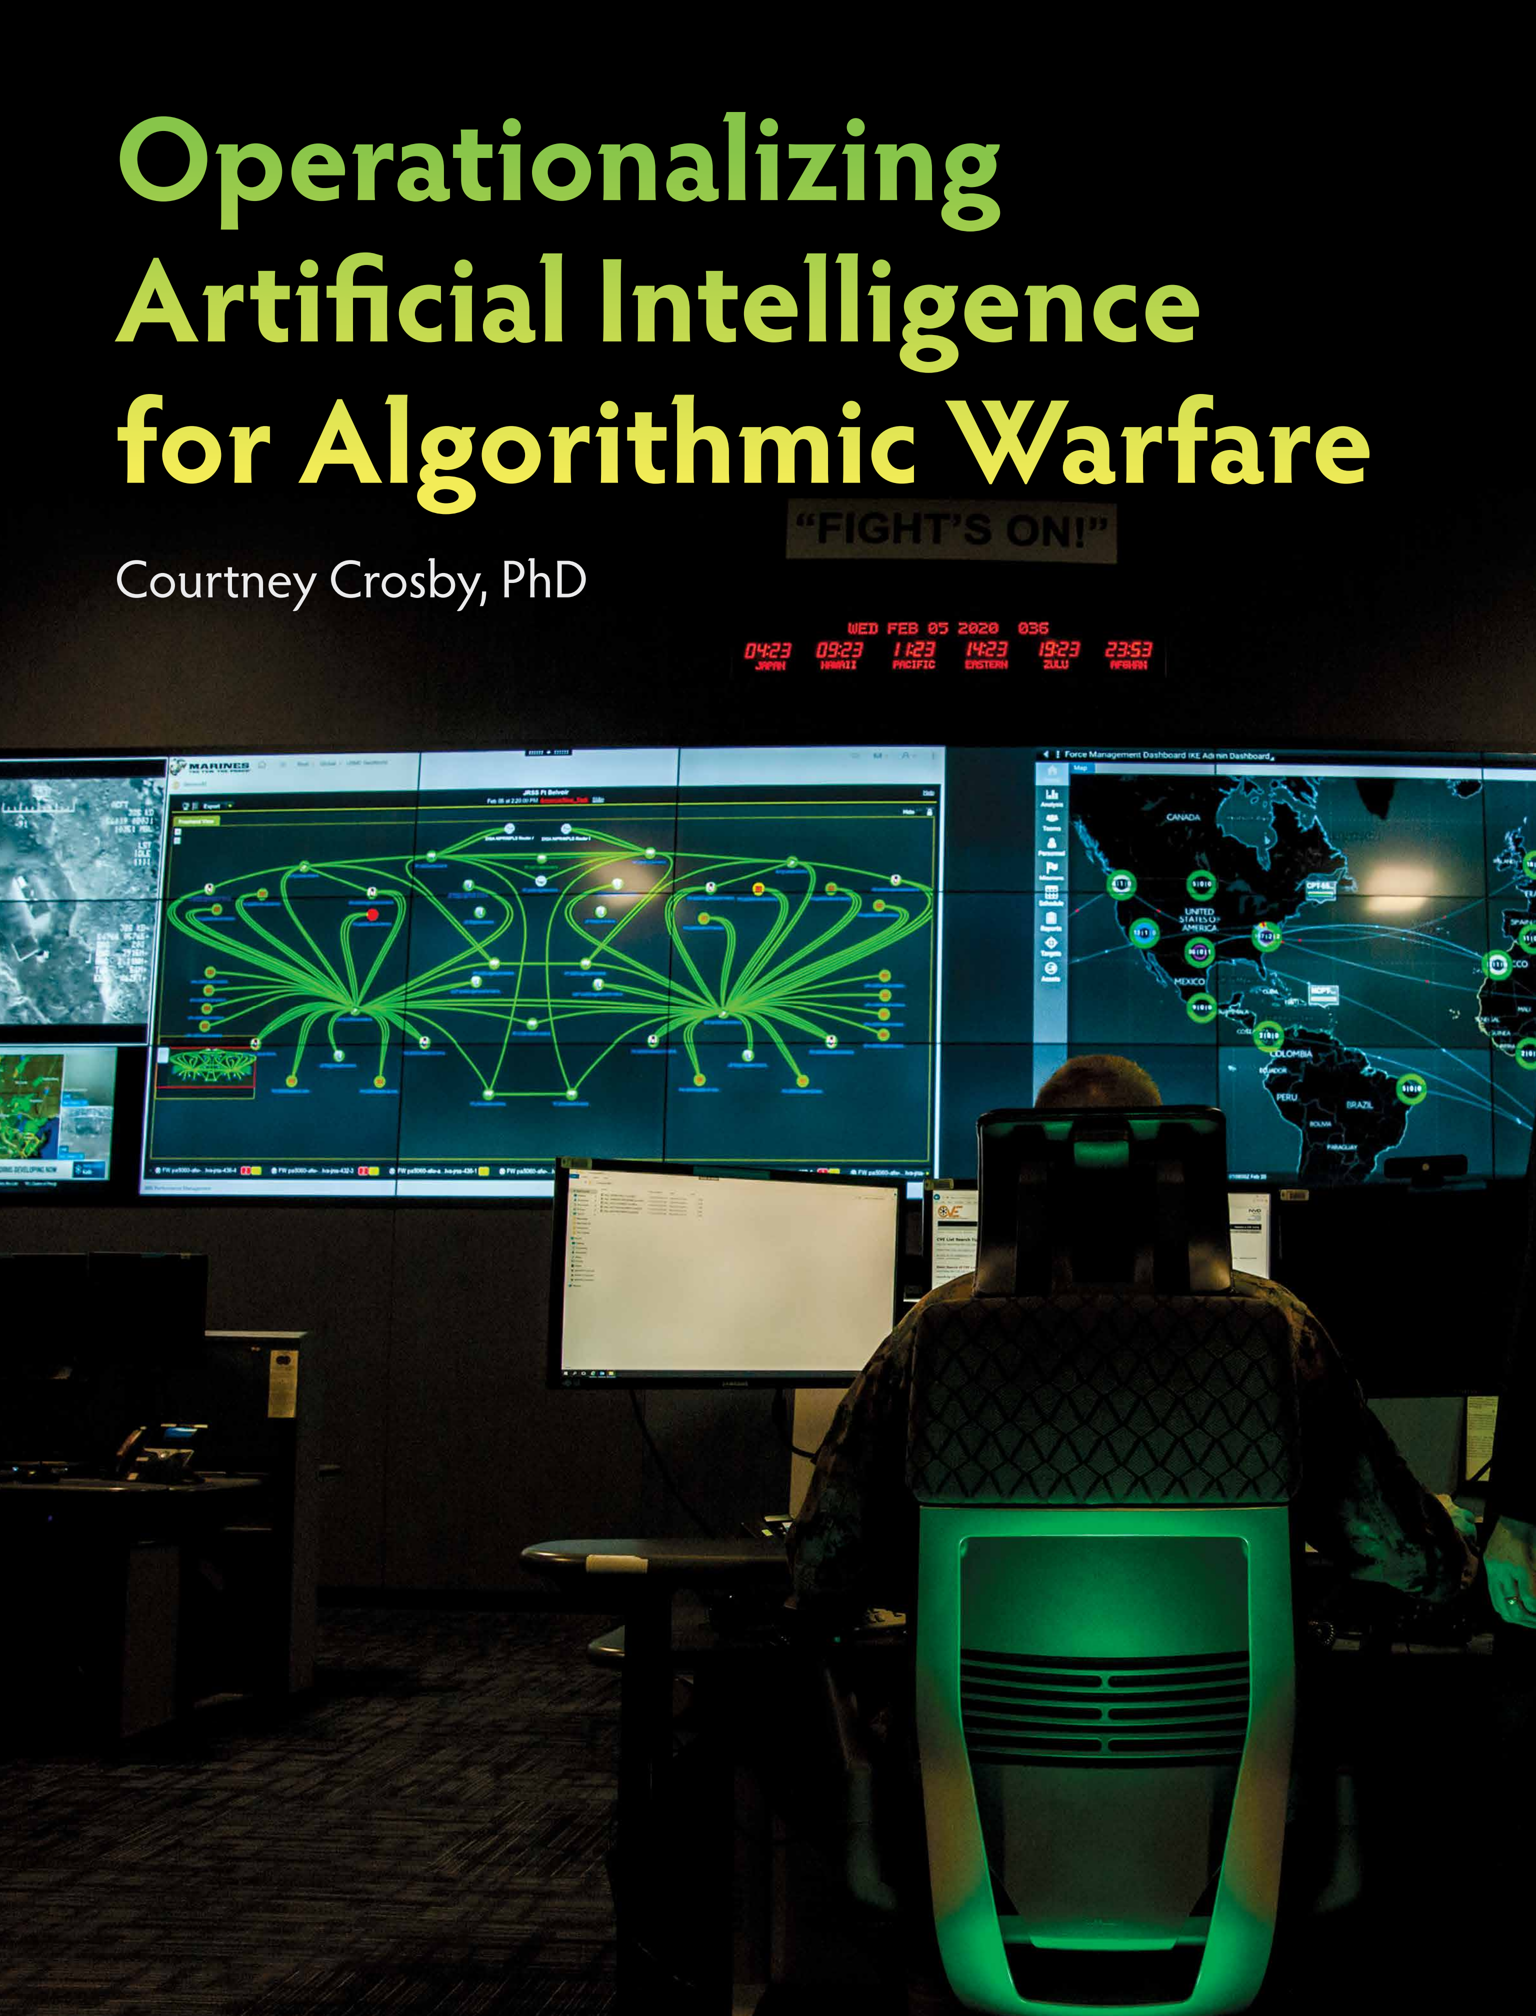

Conflict can be won or lost based on military offsets, or means with which defense units can asymmetrically combat adversarial advantages. With great-power competition, adversarial technology overmatch, and ever-expanding theaters, conventional offsets are often augmented by artificial intelligence (AI).[1] Yet, the Department of Defense's (DOD) ability to operationalize AI is nascent.[2] Initial AI programs adopted by the Pentagon focus on the transfer of commercial capabilities to the defense sector, thus highlighting technical performance and deemphasizing mission-oriented function.[3] As a result, initial pilot projects have failed to move into real-world operational environments (OE).

Operationalizing Artificial Intelligence

Operationalization hinges on the understanding that AI is not an end state but rather one way of achieving a military advantage. To that end, the technical execution of AI-related methodologies must be married to the OE. This consideration diverges from traditional thought because AI solutions are typically developed to achieve a certain statistical threshold (e.g., recall, precision), rather than a military objective (e.g., increased standoff distance).[4]

This dynamic is confounded by the term "algorithmic warfare," which currently conflates technical and military characterizations. Algorithmic warfare intends to reduce the number of warfighters in harm's way, increase decision speed in time-critical operations, and operate when and where humans are unable to operate.[5] Yet, none of those objectives speak to mathematics or computer science; they are grounded squarely in military end states. The problem is that the bridge between science, technology, engineering, and mathematics disciplines and military end states was never established before the Pentagon embarked on its AI trajectory.

The desired bridge is a framework for guiding and assessing AI operationalization, with algorithm performance on one side and mission utility on the other. Such a combination ensures that mathematical equations can prove or numerically validate an AI system while qualitative benchmarks guarantee practical application. The result is algorithmic warfare based not just on statistics but a broader architecture for operational relevancy. That relevancy is couched in five requirements:

- minimum viability,
- the ability to adapt to unknown and unknowable scenarios,
- the prioritization of insight over information,
- the requisite level of autonomy for the application, and
- battlefield readiness.

For the first time, such requirements lay the foundation for assessing military AI programs and defining success.

Marines with Marine Corps Forces Cyberspace Command observe computer operations 5 February 2020 in the cyber operations center at Lasswell Hall, Fort Meade, Maryland. Marines conduct offensive and defensive cyber operations in support of U.S. Cyber Command and operate, secure, and defend the Marine Corps Enterprise Network. (Original photo by Staff Sgt. Jacob Osborne, U.S. Marines. Photo has been modified.)

Marrying Technical Methodologies and Defense Doctrine

Developing measures of effectiveness (MOE) for military AI programs necessitates mapping research and technical methodologies (e.g., grounded theory) to DOD doctrine.[6] Without that mapping, algorithmic warfare is reduced to the process of algorithm development rather than operational deployment. For example, a computer vision algorithm designed to detect objects in a video (e.g., geospatial intelligence analysis) is reduced to the number of vehicles the model finds or how accurately it finds those vehicles. Success, then, is something to the effect of *the algorithm correctly finds vehicles 85 percent of the time*.

But what use is detecting vehicles 85 percent of the time to a military campaign? This is where preserving doctrinal integrity introduces context. Taking the example from above, the same algorithm is assessed not for how frequently it detects vehicles correctly but rather its impact to the mission: *analysts identify a vehicle of interest 95 percent faster* because of the model. Such an approach associates how well the algorithm was designed with its mission deployment. While this seems like common sense, and the relationship may even be represented ambiguously in project documentation, there is no single standard for one representation anywhere in the DOD.

Assessment criteria still need to remain solution independent (i.e., the criteria apply regardless of the type of intelligence, algorithm used, operational environment deployed to, or mission requirements). Thus, for this research, AI principles were codified into quantifiable properties and indicators that were system and program agnostic. Assessment criteria were also couched in a go-no-go fashion to create a logical, top-down hierarchy synonymous with relevant joint publications. The result is a baseline for regulating, monitoring, and evaluating DOD AI systems.

Courtney Crosby, PhD, is the chief scientist with Booz Allen Hamilton's Strategic Innovation Group. She has a PhD in decision sciences, with a focus on battlefield artificial intelligence (AI) systems for the Department of Defense. Crosby has deployed several times in support of various overseas missions. In her current role, she spearheads AI operationalization and innovation within the military C5ISR domain. Her previous experience includes research and development of emerging capabilities and social science analysis relevant to human machine teaming.

A Framework to Operationalize Defense Artificial Intelligence

As previously stated, operationalized AI is AI defined by five aspects of mission utility: minimum viability, the ability to adapt to unknown and unknowable scenarios, the prioritization of insight over information, the requisite level of autonomy for the application, and battlefield readiness. Each of these MOEs is fundamental to algorithmic warfare.[7] Analysis of this information results in a comprehensive framework of indicators and effects for each of those MOEs. The entire framework is underpinned by doctrinal definitions and procedures.

Measuring Effectiveness

The military process for measuring effectiveness relies on a go-no-go, top-down architecture. This means that a measure exists only if every single *indicator* of that measure also exists. Similarly, an indicator is present only if all *effects* of that indicator are also present.[8] It is a binary, all-or-nothing process that can be applied to AI as readily as conventional military activity.

In the conventional case of high-value target (HVT) pattern-of-life analysis, an MOE would define *one* desired result of a military campaign (e.g., the HVT moves out of the area of responsibility [AOR]). All defined indicators of that MOE must be met so that success cannot be called arbitrarily or selectively. For example, intelligence should indicate that (a) the HVT is detected in a new AOR, (b) known HVT associates are detected in the new AOR, and (c) the HVT acquires basic life support systems (e.g., housing, transportation) in the new AOR. Subsequent effects follow the same process: effects that support indicator "a" may include identification of known physical signatures and detection of communication signals.

So, while conventional and AI MOEs differ in their tactical execution, the underlying system for decision-making validation is the same. AI MOEs can only be validated if there is a baseline understanding of the AI domain, much in the same way that MOEs developed by the intel branch could not be validated by combat arms.

A display demonstrates a vehicle and person recognition system for law enforcement 1 November 2017 during the NVIDIA GPU Technology Conference in Washington, D.C. The conference showcased artificial intelligence, deep learning, virtual reality, and autonomous machines. (Photo by Saul Loeb, Agence France-Presse)

Describing Effectiveness–a Technical Wave Top

Algorithmic warfare is warfare conducted through artificially intelligent means. Artificially intelligent means are those that are not only intelligent (collecting and applying insight) but also artificial (acting on intelligence in a way that humans cannot). Without human intervention, systems must learn how to represent data for themselves.[9] Another term for this is called *machine learning*. There are different types of machine learning, but when it comes to the battlefield, *unsupervised* machine learning will become the gold standard due to its flexibility and capacity to derive outputs from unknown and unstructured information.[10] Within this gold standard, a specific methodology called *deep learning* is unique in its ability to represent complex problems more precisely.[11] Given the dynamic nature of the battlefield, the ability to represent complex problems more precisely is paramount.

Thus, algorithmic warfare can only be enabled by (a) working systems (minimally viable) capable of (b) learning on their own from unknown and unknowable scenarios (unsupervised) while (c) converting a complex battlefield environment into a useful insight (deep-learning enabled) (d) with little to no guidance

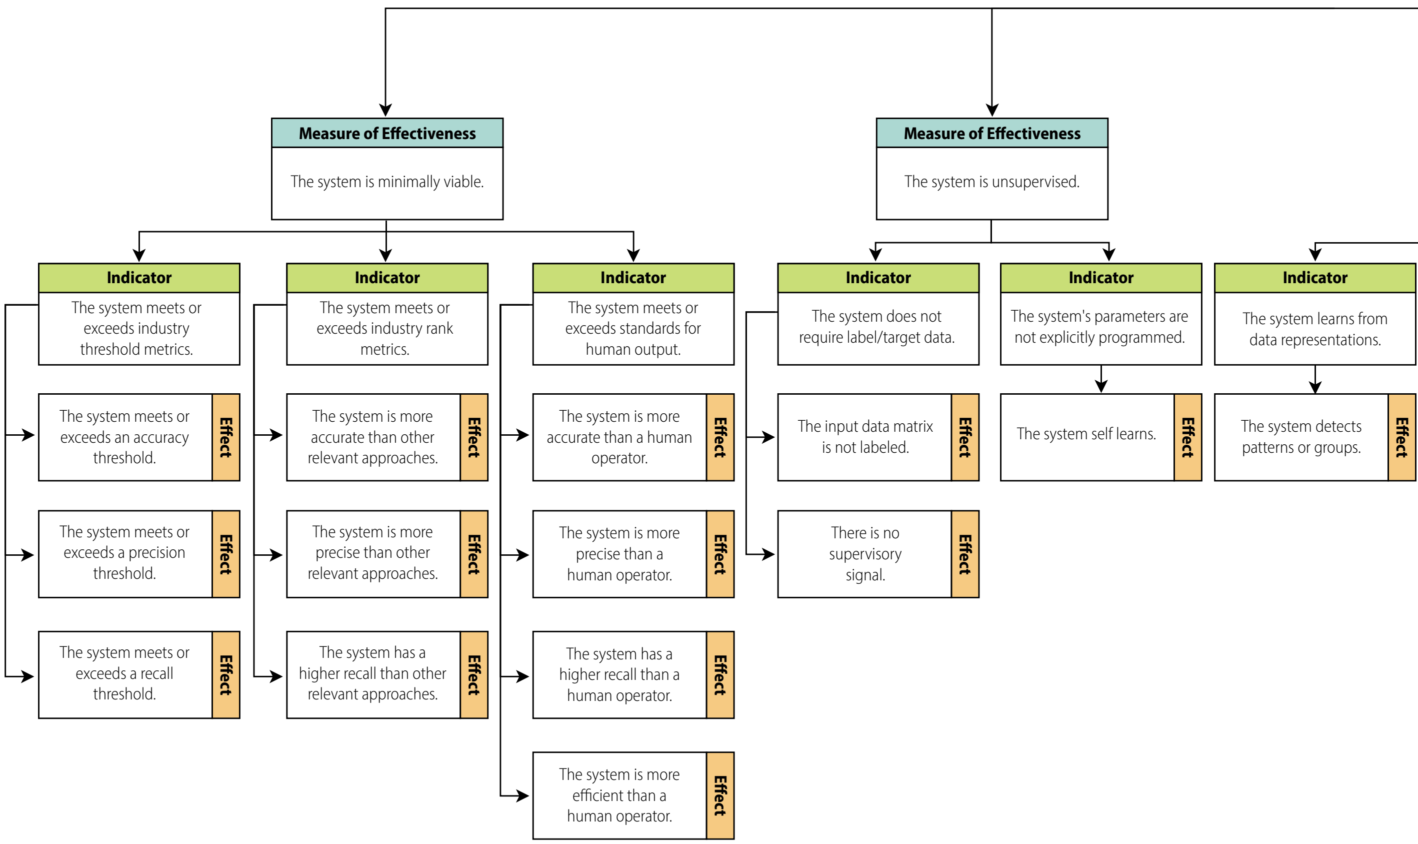

Figure. Measures of Effectiveness for Algorithmic (Artificially Intelligent) Warfare

(autonomous) and (e) in a live mission environment (battlefield ready). These MOEs and the architecture in the figure are the first steps in operationalizing AI; they lay the groundwork for how to coalesce technical and operational factors while also standardizing "success" across any AI program.

Operational Artificial Intelligence has to Work

Minimum viability tests whether algorithmic warfare positively changes the operational environment. "Positively changing the OE" means that there exists a competitive advantage and performance improvement justifying AI deployment. That justification comes from industry metrics (technical factors), ranking against similar systems, and utility to the human operator.

In the example of translation, a natural language processing algorithm would be minimally viable if (1) industry metrics confirmed that it accurately translated ground truth data from and into the correct languages, (2) the algorithm outperformed other available algorithms in the same technical class and OE, and (3) the machine translation outperformed a human.

The competitive advantage and performance improvement factors associated with minimum viability are necessary because without them, nonalgorithmically derived warfare would be more effective—thus, negating the need for operationalized AI.

Flexible and Adaptable Systems

Remember that unsupervised algorithms are ideal for live missions due to their flexibility and ability to derive insight even in unknown scenarios.[12] In short, unsupervised

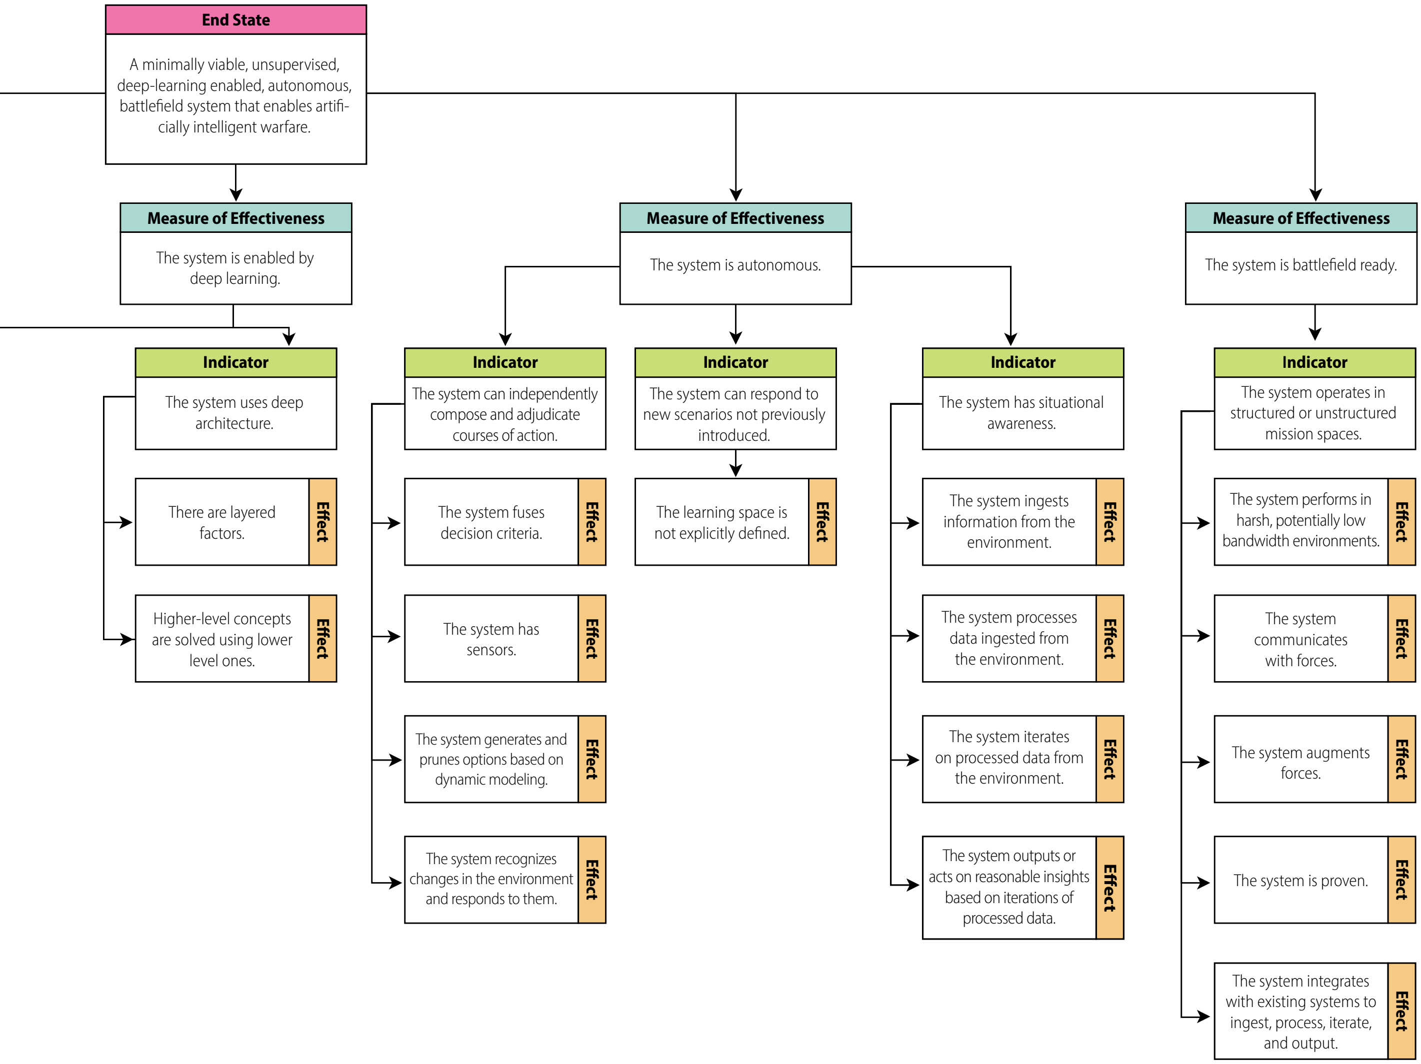

(Figure by author)

systems can operate without predetermined information and learn as new information becomes available.

A conventional equivalent can be drawn from an enemy engagement example. For instance, deployed service members do not know how a firefight will unfold until after it is over. Yet, they are expected to respond appropriately to enemy fire without warning and draw relevant conclusions about novel adversarial movement and activity.

Successful algorithmic warfare programs will need to exhibit the same adaptability of service members in their tactical execution and ability to learn over time.

Reducing Mission Complexity

Recall that deep learning reduces complexity.[13] Complexity reduction in a live mission is about how information is represented and understood. Just as with humans, effective algorithmic warfare is predicated on pattern detection, reasoning, and problem-solving.

Pattern detection is essentially acquiring knowledge that can then be generalized to predict future, unknown scenarios. Suppose that a nonaviation-branch service member deployed to an airfield sees a helicopter fly overhead. That person notices the helicopter's unique physical features, such as the overall size or a tandem rotor. The unique features differentiate the helicopter from other variations, and over time, the service member can down select the correct helicopter within an entire fleet using the learned visual cues. AI recognizes visual patterns much in the same way; helicopter characteristics are learned repetitively with subsequent sightings. Then those characteristics are generalized to differentiate one helicopter from another or a helicopter from a nonhelicopter.

Reasoning refines that knowledge acquisition in order to detect subtleties in the environment and to logically associate those subtleties. For example, if helicopters are never seen with certain weather patterns, reasoning would deduce that weather (a secondary element of the OE) influences flyability. With AI, poor weather would add secondary confirmation that a flying object without a rotor was not a helicopter.

Finally, sequential problem-solving breaks a large problem (i.e., how to fly a helicopter) into smaller problems (i.e., what is the flight path, how much fuel is available, how many pilots are needed, etc.). Thus, without complexity reduction, algorithm warfare would lack the ability to convert information to insight.

Operating with Little to No Guidance

Since algorithmic warfare assumes that other-than-human means are leveraged, AI must independently compose and adjudicate courses of action. And AI has to complete that adjudication based on its own decision-making, responsiveness, and situational awareness.

Decision-making is a matter of developing and resolving choices within the environment. In a convention setting, a commander faced with conflicting intelligence, surveillance, and reconnaissance flight paths would develop an asset prioritization matrix and then deconflict based on those requirements. This is not a matter exclusively of producing viable options but also figuring out which of those options is most beneficial to the overall mission. In order to do that, the system must be able to fuse decision criteria (e.g., number of assets, collection requirements, flight times, etc.). Sensors must be present to define decision criteria (e.g., aircraft fuel gauges or human/verbal cues). Then, all available options have to be pruned. Finally, the system has to recognize changes in the current state and respond to new information generated by that change (i.e., an aerial asset's time on station is ending so deconfliction is no longer needed).

Responsiveness complements decisiveness. That is, can the system respond appropriately to a scenario it has never seen before on the timeline required? To do so, the system has to have the requisite functions for situational awareness: ingestion, processing, iteration, and action. All indicators together ensure that operationalized AI improves decision timelines, not inhibits them.

> “Since mission constraints are vast, artificial intelligence cannot be developed in a laboratory without forethought on how it will operate in the real world.”

Moving Artificial Intelligence into the Real World

Battlefield readiness is a measure of whether the system can function in live mission spaces. Since mission constraints are vast, AI cannot be developed in a laboratory without forethought on how it will operate in the real world. To be clear, the limitations of laboratory AI are not circumvented by the battlefield; they are amplified. Open architectures are restricted by military infrastructure. Agnostic pipelines are bogged by siloed, legacy systems. Pervasive, high-speed networking becomes sporadic or intermittent once deployed forward. And the uncleared experts universal to the commercial sector are replaced by access-limited user communities with little to no AI expertise.

In short, AI must complement, rather than confuse, ongoing operations. Addressing mission constraints from the onset must then include integration and communication with existing systems. Additionally, that integration should be tested or qualified so utility, and the left/right limits of that utility, is proven prior to deployment. This would occur much in the same way that military personnel are range qualified for deployability, or conversely, how poor fitness testing can result in nondeployability.

Together, the five MOEs for operationalized AI represent standard thresholds for initial and full operating

One objective of the development of military artificial intelligence is to network soldiers directly with unmanned vehicles on the battlefield in human-intelligent agent teams that will speed the collection of intelligence, identification of targets, and execution of fire missions. (Illustration courtesy of the U.S. Army)

capabilities (IOC/FOC). IOC/FOC determinations made using the decision gates in the MOE framework will accelerate AI adoption and improve the United States' positioning in the algorithmic warfare domain.

Recommendations

Without a framework for operationalizing AI in support of algorithmic warfare, current DOD programs will fail. The framework presented in this article is the first to define success within the defense AI space and will provide necessary accountability measures for government oversight.

While the intent of this article is an agnostic solution to algorithmic warfare, additional research is necessary. Funding should be earmarked for cascading this framework to specific systems, disciplines, and programs. In support of that effort, access to both classified materials and quantitative experimentation of classified systems will be critical. Quantitative experimentation would not only serve to validate the premise of this article but also begin creating a network to compare and improve defense AI testing and evaluation. That is, continued, consistent use of the MOE architecture across multiple environments, systems, and problem sets would align AI projects under a single, common assessment framework. To that end, the MOE architecture presented in this article supports two functions: (1) to realize a more effective system by iteratively improving go-no-go decision gate results and (2) to decide between various systems by comparing respective MOEs.

Strategically, the architecture outlined in the figure (on pages 46–47) should be integrated into DOD acquisition, technology, and logistics processes. Current paradigms are not built for the exponential growth and nontraditional nature of AI programs. Calibrating current and future DOD AI solutions around prevailing evaluation criteria will enable standardization while speeding up time-consuming acquisition processes. Further, organizations responsible for enterprise AI activities should standardize the framework across their efforts for more rapid transition of applied research and development into operational use.

Organizational efforts should not stop at policy though. Currently, the DOD has no mechanism for leveraging military personnel for AI activities. Specifically, there is no military occupational specialty (MOS) related to artificial intelligence and also no official system for identifying and assigning skilled personnel to AI programs. The result is a lack of available hybrid talent; that is, personnel versed in both AI and the mission. Standing up a data science or AI-oriented MOS, similar to what occurred in the cyber domain, would make the operationalization of AI capabilities more sustainable. It would also augment the small pool of cleared AI professionals with an increasing number of qualified military personnel. Alternatively, the traditional MOS could adapt to the modern characteristics of warfare. For example, discipline-specific intelligence analysts may not be relevant in a world where multi-intelligence fusion is pervasive. Modifying or adding AI skills identifiers or specializations would curb MOS relevancy decline.

Tactically, the Pentagon's push for AI needs to be accompanied by a ground-up movement so that adopting organizations are not simply handed a capability without context. Instead, they should have an active voice in the offsets they bring to the fight. Grassroots efforts may include conducting impact analyses and stress tests at the unit level prior to IOC/FOC design plans to understand vulnerabilities and prioritize requirements.

Conclusion

Operationalizing AI is an inherently mission-centric endeavor that must make sense tactically for there to be any strategic impact. Until there is tangible return on investment for units on the ground, widespread hesitation around the value of algorithmic warfare will persist; as a result, adversarial overmatch will become an increasingly unwinnable reality.

The DOD cannot continue to execute AI programs without a framework for operationalizing those programs.[14] The architecture presented in this article does just that by accelerating and standardizing the government's efforts to develop AI capabilities through highly inventive, operationally appealing technology.[15] ■

Notes

1. U.S. Deputy Secretary of Defense Memorandum, "Establishment of an Algorithmic Warfare Cross-Functional Team (Project Maven)," 26 April 2017, accessed 4 March 2020, https://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_aw-cft_project_maven.pdf; Andrew Ilachinski, "AI, Robots, and Swarms: Issues, Questions, and Recommended Studies" (Arlington, VA: CNA, January 2017), accessed 4 March 2020, https://www.cna.org/CNA_files/PDF/DRM-2017-U-014796-Final.pdf.

2. Defense Science Board, *Report of the Defense Science Board Summer Study on Autonomy* (Washington, DC: Department of Defense, June 2016), table 1, accessed 11 March 2020, https://www.hsdl.org/?view&did=794641; *Unmanned Systems Integrated Roadmap, FY2013–2038* (Washington, DC: Department of Defense, 2013), accessed 5 March 2020, https://archive.defense.gov/pubs/DOD-USRM-2013.pdf.

3. Jesse Ellman, Lisa Samp, and Gabriel Coll, *Assessing the Third Offset Strategy* (Washington, DC: Center for Strategic and International Studies, March 2017), 6–8, accessed 5 March 2020, https://csis-prod.s3.amazonaws.com/s3fs-public/publication/170302_Ellman_ThirdOffsetStrategySummary_Web.pdf; Peter Stone et al., *Artificial Intelligence and Life in 2030: One Hundred Year Study on Artificial Intelligence* (Stanford, CA: Stanford University, September 2016), accessed 5 March 2020, http://ai100.stanford.edu/2016-report; U.S. Army Training and Doctrine Command (TRADOC) Pamphlet (TP) 525-3-1, *The U.S. Army Operating Concept: Win in a Complex World, 2020–2040* (Fort Eustis, VA: TRADOC, 2014 [obsolete]), superseded by TP 525-3-1, *The U.S. Army in Multi-Domain Operations 2028* (Fort Eustis, VA: TRADOC, 6 December 2018), accessed 5 March 2020, http://www.tradoc.army.mil/tpubs/pams/tp525-3-1.pdf.

4. *Advancing the Science and Acceptance of Autonomy for Future Defense Systems: Hearing Before the Subcommittee on Emerging Threats and Capabilities of the Comm. on Armed Services*, 114th Cong. 1 (2015), accessed 5 March 2020, https://www.hsdl.org/?view&did=793840; U.S. Army Capabilities Integration Center, *Robotic and Autonomous Systems Strategy* (Fort Eustis, VA: TRADOC, March 2017), accessed 5 March 2020, https://www.tradoc.army.mil/Portals/14/Documents/RAS_Strategy.pdf; Defense Science Board, *Report of the Defense Science Board Summer Study on Autonomy*; Executive Office of the President, *Preparing for the Future of Artificial Intelligence* (Washington, DC: Office of Science and Technology, October 2016), accessed 5 March 2020, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf.

5. Defense Science Board, *Report of the Defense Science Board Summer Study on Autonomy*.

6. Yoshua Bengio, Aaron Courville, and Pascal Vincent, "Representation Learning: A Review and New Perspectives" (Ithaca, NY: Cornell University, 23 April 2014), accessed 5 March 2020, https://arxiv.org/pdf/1206.5538.pdf; Li Deng, "Three Classes of Deep Learning Architectures and Their Applications: A Tutorial Survey"

(New York: Asia-Pacific Signal and Information Processing Association Transactions on Signal and Information Processing, January 2012), accessed 5 March 2020, https://pdfs.semanticscholar.org/5bd4/177440c17dad736f1e0d2227694d612f5a59.pdf; Francesco Gagliardi, "The Necessity of Machine Learning and Epistemology in the Development of Categorization Theories: A Case Study in Prototype-Exemplar Debate," in *AI*IA 2009: Emergent Perspectives in Artificial Intelligence*, Lecture Notes in Computer Science series, ed. Roberto Serra and Rita Cucchiara, vol. 5883 (Berlin: Springer, 2009), 182–91; Geoffrey E. Hinton et al., "Improving Neural Networks by Preventing Co-adaptation of Feature Detectors" (Ithaca, NY: Cornell University, 2012, accessed 5 March 2020, https://arxiv.org/abs/1207.0580; Nikolaus Kriegeskorte, "Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing," *Annual Review of Vision Science* 1 (2015): 417–46; Rufin VanRullen, "Perception Science in the Age of Deep Neural Networks," *Frontiers in Psychology* 8 (2 February 2017); Melanie Birks and Jane Mills, *Grounded Theory: A Practical Guide*, 2nd ed. (New York: Sage, 2015); Kathy Charmaz, *Constructing Grounded Theory* (London: Sage, 2014), accessed 5 March 2020, http://www.sxf.uevora.pt/wp-content/uploads/2013/03/Charmaz_2006.pdf.

7. Chairman of the Joint Chiefs of Staff, *The National Military Strategy of the United States of America, 2015* (Washington, DC: Joint Chiefs of Staff, 2015), accessed 5 March 2020, http://www.jcs.mil/Portals/36/Documents/Publications/2015_National_Military_Strategy.pdf.

8. TP 525-3-1, *The U.S. Army Operating Concept*, 21; Richard K. Bullock, "Theory of Effectiveness Measurement" (PhD diss., Air Force Institute of Technology, 2006), 58, accessed 5 March 2020, http://www.dtic.mil/dtic/tr/fulltext/u2/a472028.pdf; Army Doctrine Reference Publication (ADRP) 1-02, *Operational Terms and Military Symbols* (Washington, DC: U.S. Government Printing Office, 2012 [obsolete]), superseded by ADRP 1-02, *Terms and Military Symbols*, 24 September 2013; Chairman of the Joint Chiefs of Staff, *The National Military Strategy of the United States of America, 2015*.

9. Arthur L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development* 3, no. 3 (1959): 210–29; Jürgen Schmidhuber, "Deep Learning in Neural Networks: An Overview," *Neural Networks* 61 (January 2015): 85–117.

10. Alain Droniou, Serena Ivaldi, and Olivier Sigaud, "Deep Unsupervised Network for Multimodal Perception, Representation and Classification," *Robotics and Autonomous Systems* 71 (September 2015): 83–98; Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar, *Foundations of Machine Learning* (Cambridge, MA: MIT Press, 2012), accessed 5 March 2020, https://mitpress.mit.edu/books/foundations-machine-learning; Kriegeskorte, "Deep Neural Networks."

11. Volodymyr Mnih et al., "Human-Level Control Through Deep Reinforcement Learning," *Nature* 518, no. 7540 (26 February 2015): 529–33; Deng, "Three Classes of Deep Learning Architectures and Their Applications"; Kriegeskorte, "Deep Neural Networks"; Schmidhuber, "Deep Learning in Neural Networks."

12. Droniou et al., "Deep Unsupervised Network"; Kriegeskorte, "Deep Neural Networks"; Mnih et al., "Human-Level Control."

13. Deng, "Three Classes of Deep Learning"; Kriegeskorte, "Deep Neural Networks"; Schmidhuber, "Deep Learning in Neural Networks."

14. Stew Magnuson, "DoD Making Big Push to Catch Up on Artificial Intelligence," *National Defense* (website), 13 June 2017, accessed 5 March 2020, http://www.nationaldefensemagazine.org/articles/2017/6/13/dod-making-big-push-to-catch-up-on-artificial-intelligence; JASON [advisory group], *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD* (McLean, VA: MITRE Corporation, January 2017), accessed 5 March 2020, https://apps.dtic.mil/dtic/tr/fulltext/u2/1024432.pdf.

15. Mark A. Milley, Memorandum for all Army leaders, "Army Readiness Guidance, Calendar Year 2016–17," 20 January 2016, accessed 5 March 2020, https://www.army.mil/e2/downloads/rv7/standto/docs/army_readiness_guidance.pdf; U.S. Army Capabilities Integration Center, *Robotic and Autonomous Systems Strategy*; Defense Science Board, "Report on the Defense of the Science Board"; TP 525-3-1, *The U.S. Army Operating Concept*; Noel Sproles, "Coming to Grips with Measures of Effectiveness," *Systems Engineering* 3, no. 1 (2000): 50–58.