Vulnerabilities and Challenges of Integrating AI into Future Air Force Intelligence Systems

Major Mariko Hart
Student 2825

ACTS 2.0 RTF

20 Apr 2020

Abstract: DoD and Air Force leaders have identified Artificial Intelligence (AI) as a game-changing technology that will help the Air Force Intelligence, Surveillance, and Reconnaissance (ISR) enterprise overcome perennial challenges with speed and scale of intelligence analysis required for great power conflict. The Sensing Grid concept (recently renamed Sensor Integration) was introduced as a future framework to integrate AI and Cognitive Modeling tools for Air Force ISR, but there is little discussion about adversary threats and ethical considerations that should inform the design and functionality of the system. In order to prepare the human and organizational terrain within the Air Force to integrate a highly automated, AI-enabled intelligence analysis system, enterprise leaders must advocate for a human-centric design that takes lessons learned from historical human-machine teaming successes and failures. Leaders must also take a proactive approach to training the Air Force ISR workforce to team effectively with revolutionary but imperfect AI technology.

Problem Statement

According to the Air Force Warfighting Integration Capability's (AFWIC) Sensing Cross Functional Team, the current state of Air Force intelligence, surveillance, and reconnaissance (ISR) is highly specialized, proprietary, and relies too heavily on human intensive reach-back processes.[1] As planners look toward great power conflict in the future, they assess present-day hardware and analytical processes will be insufficient to establish decision advantage against a peer adversary, with intelligence lagging behind in the speed and scale required for victory.[2] The Air Force A2's *Next Generation ISR Dominance Flight Plan* is similarly critical of the current ISR enterprise, advocating for a departure from today's "industrial-age, single-domain approach" to pursue "architecture and infrastructure to enable machine intelligence, including automation, human-machine teaming, and ultimately, artificial intelligence."[3] While providing Airmen with faster, smarter tools to craft and share assessments is a priority for Air Force senior leaders, introducing greater levels of automation and machine-led sensemaking presents a host of new concerns for the intelligence community. Is it safe to rely on algorithms for shortcuts, considering the threats of tampering and deliberate misinformation these tools could encounter? Is it ethical to pursue warfare enabled by automated weapon systems at all? If so, what risks does the intelligence community incur by employing automated tools to produce critical intelligence assessments with greater speed?

The term Artificial Intelligence (AI) is defined by the Department of Defense's Joint Artificial Intelligence Center as "the ability of machines to perform tasks that normally require human intelligence – for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action."[4] Seniors leaders are hopeful that AI can soon power software suites human analysts use to make assessments and enable physical systems in

more autonomous applications.[5] Machine Learning (ML) is defined by the Defense Advanced Research Projects Agency (DARPA) as a field within AI "that applies statistical and probabilistic methods to large data sets" and can apply derived models to future data samples.[6] A popular approach to harnessing the benefits of ML is through Deep Neural Networks (DNNs), which can be trained to perform a range of classification and prediction tasks using historical data.[7] Although it is not specifically referenced in AFWIC or A2's publications, using AI, ML, and DNNs in applications simulating human thought processes is a hybrid field of computer science and psychology called *cognitive modeling*. AI, ML, DNNs and cognitive modeling concepts are a pivotal part of the transformation to a digital, network-centric approach to intelligence in AFWIC's vision of the future Air Force ISR enterprise.[8]

To provide a framework for modernization initiatives in the Air Force ISR enterprise, AFWIC built the Sensing Grid concept, defined as "an assemblage of sensors, platforms, people, devices, content, and services that delivers a holistic, accurate, predictive, and timely characterization of the operating environment to decision makers."[9] Architects of the concept envision a system featuring predictive analysis tools, autonomous sensing and response, the ability to fuse multiple data sources, and edge processing, all of which are enabled by leveraging AI, ML, DNNs, data analytics, and other cognitive modeling methods.[10] Although there is no published date for the debut of the Sensing Grid, most discussions indicate the optimized system of systems is at least ten years away. Meanwhile, there is great urgency from DoD leadership to catch up to investments China and Russia have made in military AI applications, encouraging rapid prototyping and experimentation to find solutions.[11] AI is routinely identified in defense forums as the answer to making data-centric intelligence tasks faster and speed up tactical

decision-making, but this is mere conjecture if the systems involved are in nascent stages of engineering and remain unproven in the national security arena.[12]

While AFWIC and the Air Force A2 focus on investment in research and development of AI-enabled sensors and tools, the dynamic of human-machine teaming required to make the Sensing Grid safe and effective is rarely discussed. In order for the Sensing Grid to be an effective system that brings value and advancement to the way the Air Force executes ISR and analysis, enterprise leaders should advocate human-centric design in the technology, train and prepare frontline analysts to team effectively with the new system, and adapt organizational practices informed by the strengths and weaknesses of AI. Air Force leaders must acknowledge the adversarial threats and ethical problems inherent in assigning a greater number of analytical tasks to AI tools, and these concerns must inform the blueprint of the Sensing Grid. This is not to suggest that ongoing software development of the system should stall, but rather that a concurrent dialogue must take place among intelligence and materiel leaders about the role of human analysts as crucial to mitigating the downsides of an increased reliance on AI. Air Force leaders must also pursue a deliberate plan to integrate Sensing Grid components into current sensing, identifying, attributing, and sharing (SIAS) activities, preparing front line analysts for tasks of "higher-level reasoning and judgment" while acknowledging that machines should augment human tasks, not replace humans entirely.[13] While it may be tempting for leaders to make force structure changes in ISR manpower in anticipation of fewer requirements thanks human-machine teaming, premature decisions like these stand to leave the force unfit to actualize the benefits of an AI-enabled intelligence enterprise.

The pages to follow will provide a literature review on vulnerabilities and ethical problems associated with AI-enabled systems for insight on what challenges may lie ahead for

building and fielding the Sensing Grid. It will also include a discussion on what factors intelligence and materiel leaders should consider before completion and fielding this game-changing system of systems. The paper will conclude with further recommendations on how to prepare the Air Force ISR battlespace for the Sensing Grid, offering ideas on the scene-setting required to posture Airmen for operations in the Digital Age.[14] By encouraging further exploration and discussion on this topic, I hope that AI will be endorsed less often as a panacea to replace human analyst labor and make data aggregation problems disappear, and will be regarded instead as a powerful tool that Air Force ISR leaders must understand completely and integrate carefully.

Literature Review

Most recent research on vulnerabilities associated with applying AI to cognitive tasks highlights the danger of adversarial examples, which modify inputs to DNNs and cause systems they control to malfunction in a variety of ways.[15] Adversarial inputs can be physical or non-physical, and can affect a variety of data classifiers sorting media including images, audio files, and text.[16] The most frequently referenced example of robust physical deception is an experiment where engineers fool the optical sensor on a self-driving car by angling stop signs differently, causing the vehicle to miss the stop.[17] Physical deception is not a new or novel scheme in defense applications, but incorporating edge processing and automation into systems like the Sensing Grid may preclude human analysts from identifying these tactics firsthand. In the non-physical realm, training algorithms to identify patterns in a manner similar to the human brain is a challenging task. Computer Vision (CV) algorithms classify imagery very differently than human analysts and are prone to misclassifying objects when just a few pixels are out of place.[18] In less straightforward circumstances, engineers cannot explain the model's errors, spurring

initiatives for *explainable AI* by organizations like DARPA.[19] In the best case scenario, adversarial inputs are recognized as outliers and overlooked by CV models with robust training samples; at worst, they could corrupt real-world inputs and digitally remove objects or activity from a sample without a human analyst's knowledge.[20] Adversarial inputs in this context could have catastrophic consequences if they caused an analyst to miss significant activity they would normally catch unassisted.

If the objective behind applying AI, ML, and DNNs to intelligence datasets is to analyze and disseminate more information at a higher speed, Natural Language Processing (NLP) is also likely to be a part of Sensing Grid architecture. NLP models are widely employed today for personal and commercial use, with tools like Siri and Amazon Alexa using voice cues to initiate other applications. NLP models can also be used for comprehension tasks on large volumes of text or other media, using derivative data answer questions.[21] This technology is likely to be extremely useful in SIAS tasks fusing multiple sources of data, but could also be susceptible to interference.[22] Adversarial inputs in NLP can introduce erroneous sentences or replace key words with their antonyms in text files, causing a model to mischaracterize datasets in instances where there is no time or capacity for manual review.[23]

As with any scenario that layers models, it is unknown whether CV and NLP models will work together as effectively as predicted, let alone detect fabricated data like Deepfakes that enter DNNs in the unclassified domain.[24] The further human analysts are removed from source data streams where they can ordinarily detect misinformation, the more susceptible SIAS could become to false inputs. Despite this concern, guidance from the Air Force A2 suggests that there are high expectations for layered models to leverage unclassified Publicly Available Information (PAI) for seamless sensor cueing, allowing the ISR enterprise to find relevant targets more

efficiently.[25] Without a robust method of detecting fabricated media within PAI samples that cue sensors, this process could be difficult to achieve securely.

The technical complexity and potential vulnerability to tampering in automated, AI-enabled systems generates discussion about whether applying this type of technology in military operations is ethical. Although the Sensing Grid is not designed to employ weapons directly, intelligence data from the system will likely inform time critical decisions about effects in multiple domains.[26] Literature on the ethics of AI/ ML is generally critical of employing systems that function autonomously with limited windows for human intervention, following a similar logic to the ethical argument against legacy automated weapons like landmines.[27] While the Sensing Grid and its precursor systems will have much higher cognitive-behavioral attributes than a pressure-plate device, some posit that humans have equally little control over black box systems that execute layers of algorithm-to-algorithm communication before presenting the human operator with a choice or conclusion.[28]

Fortunately, AI systems may also be able to compensate in scenarios where humans are prone to ethical missteps, since machines do not experience emotions like fear or panic that can trigger risky decisions or LOAC violations.[29] Though it would be prudent to leverage this potential strength in human teaming with cognitive models, guidance from the DoD places specific value on speed as the most useful contribution of AI, which introduces more ethical quandaries. Tests of individual decision making suggest that humans are already poor evaluators of risk in complex environments, and the introduction of AI shortcuts that marginalize the value of human judgment will only lead to faster, higher-stakes conclusions.[30] When incorrect assessments or hasty decisions enabled by AI lead to catastrophic error, accountability is also a confusing ethical matter that the DoD leaders must prepare to address.[31]

The implied solution to mitigate threats of adversary tampering and ethical missteps in most literature is an optimal division of labor between the human controller and autonomous, AI-enabled system. Unsurprisingly, there are many perspectives on what this should look like and how it might apply to a system like the Sensing Grid. Some suggest there is no room in the framework of international agreements for employment of automated weapons systems, comparing their lack culpability to child soldiers.[32] Others believe AI tools will not result in an unacceptable loss of control if processes like Joint Targeting carry on with the same rigor and rules of engagement.[33] While there is perceived urgency to deliver Sensing Grid capabilities to the joint intelligence community through the purchase of existing commercial software, if DoD leaders wish to reduce the previously discussed risks, engineers, requirements owners, and analysts must commit to thoughtful discussion on where AI applications will be most helpful in the ISR enterprise and where they have potential to do harm.[34]

Discussion of Findings

When it comes to investing in programs to build the future ISR enterprise powered by AI and cognitive modeling applications, the DoD and Air Force provide limited direction outside of the need to invest quickly and partner with universities and national laboratories.[35] There is limited guidance on what the department desires in an AI investment outside of systems that have "a lower risk of accidents; are more resilient…to hacking and adversarial spoofing [and] demonstrate less unexpected behavior."[36] This list of attributes provides valid engineering concerns about robustness and security, but omits any discussion of the desired function or teaming construct where AI should be employed. The lack of specificity may be a symptom of how recently AI has debuted in DoD strategy, but complacency and contentment with investing for investment's sake is not far ahead of this scenario. It is incumbent on the community using

the technology to decide what type of partnership with cognitive models would provide the greatest benefit, yet strategic guidance appears to surrender responsibility to laboratories and industry partners, charging outsiders to identify both the problem and the solution that AI will solve. If Air Force ISR enterprise leaders do not play an active role in discussing how best to partner human analysts with AI tools, how will they evaluate whether developers are making sufficient progress with funding provided? How can the DoD trust that a solution developed by non-operational partners adequately addresses safety and ethics concerns? At what point does AI transition from a tenuous research project to a tenable solution improving the speed and accuracy of SIAS?

A more productive approach to discussing AI and its prospective functions in intelligence is to refrain from treating it like a magic bullet that is too ill-defined to examine at all. While applying cognitive modeling to intelligence processes may be new, technology enabling automation in warfighting has existed for decades.[37] It is imperative that leaders consider the human-machine teaming constructs that already exist in modern warfare to derive lessons for both design and integration of the Sensing Grid. For the Air Force ISR enterprise, analyzing current and historical teaming of human analysts, airborne sensors, and theater decisionmakers is a useful exercise.[38] Measures of performance for airborne ISR sensors are normally evaluated through factors like responsiveness and accuracy of sensor output, but understanding the analytical and decision-making processes that sensor data triggers is also important. Spectral imaging sensors, for example, can be employed as *anomaly* detectors highlighting unusual objects or activity for a human analyst to review and report.[39] The report could be disseminated to an operational leader, who then makes a decision based on the intelligence to order an airstrike on the source of the anomalous activity. If this chain of events occurs habitually over the course

of an operation, the interaction between sensor and humans in the loop may start to change, whereas the sensor is subconsciously reclassified as a *threat* detector.[40] The performance specifications of the sensor in this scenario did not change, but the humans in the teaming relationship began to apply a different value to the sensor output over time, which could be the influence of external motivators. While most analysts know that it is incorrect and dangerous to assume that all anomalies are threats, the subtle ways in which human-machine teaming relationships can evolve to distort human judgement is cause for concern. In order to ensure that human-machine teaming is conducted ethically, leaders must reflect on how teaming constructs could unintentionally inhibit the organization's values.[41] Demanding accuracy and robustness in new warfighting technology is reasonable, but understanding the organizational behaviors and habits the technology instigates is paramount to effective and ethical employment.

In addition to applying lessons from existing human-machine teaming within the ISR enterprise, an AI-enabled Sensing Grid should also be human-centric in its design. While this may seem obvious when building a system employed by human analysts, human factors engineering and human-machine teaming considerations are frequently a low priority in complex systems engineering projects.[42] This is due in part to traditional organizational barriers that place software engineers and human factors specialists in different departments, especially as the latter specialize in disciplines like cognitive psychology, neuroscience, and robotics that may play a limited role in some programs.[43] The consequences of failing to properly integrate human factors in complex systems can be dire as seen in Boeing's 737 Max aircraft, which was involved in two fatal accidents in 2018 and 2019.[44] Both accident reports cited the highly automated Maneuvering Characteristics Augmentation System (MCAS) software as a significant contributing factor in the aircraft mishaps.[45] While MCAS was designed to use sensor inputs to

assist with flight safety, poor human factors considerations made the system too difficult for pilots to override once automated processes were triggered.[46] Although training users to team with a new system is a natural part of onboarding, a steep learning curve caused by a lack of human factors engineering is a risk that could be mitigated by modeling both human and machine behaviors as they relate to the tasks at hand.[47] Modeling in this instance would help system architects identify the communication gaps that cause misunderstanding in a specific teaming relationship, perhaps delivering insight on how a machine can sufficiently disclose its limitations to the human operator prior to an emergency scenario.[48]

As we speculate how best to facilitate human-machine communication that adequately addresses the safety and ethics concerns associated with AI and automation, seeking consultation from visual analytics experts could provide valuable design insight.[49] Visual analytics is a scientific field that seeks to increase transparency in automated, high-volume data processing by increasing human-machine discourse through interactive visualizations.[50] Providing analysts with a teaming construct that allows them to choose how to visualize a dataset could strike a favorable balance between automated, machine-assisted data streamlining and human judgement.[51] In the best case application of visual analytics to the Sensing Grid, an analyst would comprehend the significance of a dataset with high confidence, thanks to the ability to *tune* the underlying analytical processes.[52] Ideally, visual analytics enable the user to leverage their subject matter expertise by pitching hypotheses and questions about the data to the system, allowing them to draw conclusions through the dialogue.[53] An approach within visual analytics known as semantic interaction could also be helpful, creating models that translate analyst dialogue with visual data into model adjustments, inferring and learning why the human partner is performing a routine task like highlighting, copying, etc.[54] Considering how relatively new the

previously detailed disciplines are, establishing well-defined test and evaluation standards would be an important step in preparing to integrate these and other teaming techniques into SIAS tasks.[55] Standards informed by the operational community would provide developers an understanding of what types of dialogue *explanations* from models are most helpful to contextualizing conclusions for intelligence analysis.[56] While there is still a great deal of basic, inter-disciplinary research to be completed within the field of visual analytics, deliberately applying this scientific discipline to Sensing Grid design would be a worthwhile investment in the near term for the ISR enterprise.

There are undoubtedly many challenges ahead for the directorates within Air Force Research Labs (AFRL) who have been working to build components of the Sensing Grid before the concept was formalized. The preponderance of engineers and developers integrating AI into intelligence architecture and software work in AFRL's Information Directorate (AFRL/RI) at Rome Labs, divided among multiple Core Technical Competency (CTC) teams. The Processing and Exploitation (PEX) CTC in particular will be deeply involved in developing the DNNs that enable the Sensing Grid, with a mission to "deliver fast sensemaking for situational awareness and adversarial insight to the AF, DoD and Intelligence Community".[57] Within the PEX CTC, programs are split by function into Characterization, Extreme Computing, and Comprehension and Projection programs, covering the range of steps from data extraction to advanced sensemaking.[58] Expertise in Human Factors engineering comes from Airman Systems (RH), an interdisciplinary directorate located two states away at Wright-Patterson Air Force Base. Down the line, programs in the PEX CTC may integrate with projects in development within other AFRL directorates like Sensors (RY) or Aerospace Systems (RQ), connecting RI's portion of SIAS with new airborne collection sensors and vehicles. For the time being, RI engineers work

with sample data streams from real- joint and national intelligence sources, incrementally addressing the computational challenge of sorting through a high volume of unstructured data.[59] Finding solutions to keep the size, weight, and power requirements of physical systems manageable is also a constant concern, especially in programs like Agile Condor seek to provide high levels of edge processing on airborne systems.[60]

As the earlier survey of literature indicates, building robustness and security into DNNs to prevent adversarial interference in ML is a concern for any network developer, and teams within RI are no exception.[61] DNNs already learn or fail to learn in unexpected ways in a lab environment, and introducing adversarial inputs that contradict human perception could set back any progress made on developing a useful tool.[62] Since DNNs are designed to learn and adapt when presented with new data, systems reliant on learning models also may not fit neatly into traditional DoD Technology Readiness Levels. If system continues to evolve with new datasets, it may prove difficult to define benchmarks of technology maturation where it would be appropriate for AFRL to transfer sustainment responsibilities to the Air Force Lifecycle Management Center (AFLCMC). While this point relates back to the importance of establishing test and evaluation standards for AI-enabled Sensing Grid components, it should also generate discussion about whether the handoff of complex systems between development and sustainment organizations is appropriate for this type of technology. Ideally, the teams with the greatest expertise on DNNs who build the models would maintain them throughout their lifecycle. A more likely and less disruptive course of action would be to build Sensing Grid components with an upgradeable chassis and form factor, allowing streamlined upgrades with replacement equipment as it becomes available.[63] A system design that allows for an iterative relationship between end user community and expert developers would also be advantageous in incorporating

the impending findings of DoD AI basic research collaborations taking place throughout the country. Considering the multitude of AI research investments at national labs, DARPA, the Massachusetts Institute of Technology, Carnegie Mellon, and other institutions, Air Force leaders should be thinking about how the department's return on investment could be integrated to improve design and functionality of the Sensing Grid as findings are published.

For DoD and Air Force leaders setting the scene for integration of the Sensing Grid in the future, there are other unique ethical challenges that should reconciled. If the Sensing Grid and its components deliver the rapid and robust sensemaking the concept promises, is it reasonable to expect that all frontline analysts using the system will understand how it works? In the event of a catastrophic error, would a junior analyst need to understand the technology in order to be held accountable for errors of suspected negligence?[64] The incorporation of edge processing into Sensing Grid designs is also an ethically controversial subject. Although automated data processing could save time in SIAS, how will analysts know if an edge computing program malfunctions, or if they are being deceived by an adversary?[65] Removing human cognitive labor from the sensor's edge could deliver data faster, but accuracy of the results could vary. Leaders who recognize these issues but delay addressing them because of pressure to field the technology faster than China or Russia should think through the reasoning behind this position. Any nation state competing in the perceived race to build an AI-enabled Command, Control, and ISR system is equally disincentivized to cut corners in areas like safety and security.[66] While China and Russia have forms of government that are fundamentally different than the US, the fact remains that both nations have hierarchical systems where accountability for errors and imprecision in national defense matters is also taken seriously.[67] In a manner similar to nuclear programs, the US government should lead the international community in sharing tradecraft for safe, well-

designed AI algorithms with competitors, ensuring that no nation spurs a conflict of misunderstanding due to bad situational awareness tools.[68] The best international arms control for AI is likely to come from being as transparent as possible on the findings of AI research and advocating responsible use of the technology.[69]

Recommendations

Although the Sensing Grid in its complete form is years away from coming to fruition, components of the final system will likely be fielded incrementally over the next decade. There is a great deal of work ahead to prepare the technology, the people, and the organization properly for this next generation of human-machine teaming. DoD and Air Force ISR enterprise leaders should not wait for the formal system debut to begin advocating for human-centric design in Sensing Grid technology, incorporating AI-informed training objectives into instruction for frontline analysts, and priming organizations to receive and team with the technology. When it comes to designing and building this complex system of systems, materiel leaders should proceed with caution when considering procurement of commercial, off-the-shelf software for faster data aggregation solutions. Without establishing test, evaluation, and security standards for how the Sensing Grid and its precursor systems should function, prematurely integrating multipurpose commercial software could introduce uncertainty and risk into the Sensing Grid's AI to AI interactions.

Furthermore, the desire to find a faster solution should not precede human factors considerations that will be vital to safe and productive human-machine teaming. DoD leaders should also take a critical look at plans to integrate edge processing across the Sensing Grid as a safety and ethics concern, and should think carefully about where it would be truly appropriate to divorce human perception from sensor outputs. Adaptive training models should also be

developed within the Sensing Grid system to ensure that analysts can learn from dialogue with the system about its limitations and strengths.[70] While training human analysts is one of the most obvious measures the ISR enterprise can take to mitigate threats from external meddling and ethical missteps, materiel leaders must also consider their roles in procuring well-designed, human-centric technology as an equally important safeguard.[71]

As the Defense Innovation Board's *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence* suggests, it is critical for junior officers, junior enlisted members, and civilians to receive training on AI as early as possible in their careers.[72] Although young analysts demonstrate high aptitude for learning digital applications and programs quickly, junior personnel also tend to trust technology in surprising ways.[73] Therefore it is imperative that these analysts continue to receive training on the fundamentals of intelligence analysis, making them adept at recognizing the errors and omissions of algorithms they will inevitably team with in the Sensing Grid. Air Force leaders took a pragmatic first step toward promoting AI and ML literacy in 2018 by launching a pilot program to identify Airmen with computer language experience, hoping to leverage those with coding expertise in various initiatives.[74] While this measure will be helpful in distinguishing analysts with higher degrees of digital proficiency, teaching the workforce how computer models function is probably a more useful skillset to prepare for human-machine teaming in the Sensing Grid.[75] The worst approach to preparing frontline analysts for Sensing Grid employment would be to rely on Just-in-Time Training to bridge gaps in workforce knowledge about the technology, thereby introducing a greater margin of error to SIAS activities.

To prepare the organization to receive and integrate the Sensing Grid, DoD and Air Force leaders must begin by managing the problematic expectation that AI-enabled software will

reduce the need for human analyst manpower in the ISR enterprise. Although systems like the Sensing Grid are designed to mimic human cognitive labor, analyst labor will still be critical for quality control and task management, not to mention as a safeguard to identify potential tampering or system malfunction within DNNs. Now is not the time to make any drastic force structure adjustments in anticipation of a technological advancement that is years away from fielding. Furthermore, there is little discussion so far on how the Sensing Grid will integrate data from the joint force, or whether Combatant Commands will be permitted to have their own distinct data strategies and intelligence resources as they do today.[76] If the Sensing Grid fails to provide analysts with greater access to sources of intelligence because of artificial seams from a service component or Geographic Combatant Command, is the system truly doing what its architects advertise? These issues must be addressed and reconciled at the Joint Staff level. Finally, organizations that leverage intelligence that comes from the Sensing Grid must be cognizant of the biases and shortcuts they are susceptible to when they team with machines. Understanding how external pressure and rules of engagement can cause a failure to question machine output is vital to improving human-machine partnerships that truly make SIAS more effective.

Conclusion

The research investments made by both the DoD and the Air Force on the employment of AI in intelligence applications will be of utmost importance in determining how the force should prepare for human-machine teaming with the Sensing Grid. Training both enterprise leaders and frontline analysts on the ethical quandaries and potential for adversary meddling in automated, AI-enabled SIAS will be crucial to protecting organizations from propagating misinformation. Luckily, it is not too late for DoD and Air Force ISR enterprise leaders to advocate for human-

centric designs and training models within the Sensing Grid system, as engineers at AFRL continue to work toward a secure, pragmatic solution for the force. Leaders must be cognizant of organizational inclinations to trade accuracy for speed, and understand that deliberate, phased integration of a well-designed system will be worth the wait.

Notes

[1] The author wishes to thank Dr. Paul Hoffman, Dr. Leslie Blaha, Dr. Christian LaBiere, and Dr. David Danks for their thoughtful comments and suggestions. All errors found herein are my own. Air Force Warfighting Integration Capability (AFWIC), "Joint All-Domain Command and Control Sensing Grid" (slideshow, 14 Nov 2019).
[2] Ibid.
[3] Lt Gen Dash Jamieson, "Next Generation ISR Dominance Flight Plan Summary," (Washington, DC, United States Air Force Deputy Chief of Staff for Intelligence, Surveillance, and Reconnaissance), https://www.af.mil/Portals/1/documents/5/isrflightplan.pdf.
[4] Department of Defense, "Summary of the 2018 Department of Defense Artificial Intelligence Strategy: Harnessing AI to Advance Our Security and Prosperity" (Washington, DC: Office of the Secretary of Defense, Feb 2019), 5.
[5] Ibid.
[6] Defense Advanced Research Projects Agency, "AI Next Campaign," DARPA.mil, accessed 16 Nov 2019, https://www.darpa.mil/work-with-us/ai-next-campaign.
[7] Ibid.
[8] Air Force Warfighting Integration Capability (AFWIC), "Joint All-Domain Command and Control Sensing Grid" (slideshow, 14 Nov 2019), 5.
[9] ***In Spring 2020, AFWIC renamed the initiative *Sensor Integration* under the JADC2 framework, although *Sensing Grid* is still used by AF/A2.** AFWIC, "Joint All-Domain Command and Control Sensing Grid," 7.
[10] Ibid.
[11] Department of Defense, Summary of the 2018 Department of Defense Artificial Intelligence Strategy: Harnessing AI to Advance Our Security and Prosperity (Washington, DC: Office of the Secretary of Defense, Feb 2019), 7.
[12] Ibid.
[13] Ibid.
[14] Jamieson, "Next Generation ISR Dominance Flight Plan Summary."
[15] Kevin Eykholt et al., "Robust Physical-World Attacks on Deep Learning Visual Classification," (paper presented at 2018 Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, June 2018), 1.
[16] Ibid. Naveed Akhtar and Ajmal Mian, "Threat of Adversarial Attacks on Deep Learning in Computer Vision: A Survey," *IEEE Access* Vol 6 (19 Feb 2018), 26. / Wei Emma Zhang et al., "Adversarial Attacks on Deep Learning Models in Natural Language Processing: A Survey" (The University of Adelaide: Australia, Oct 2019), 2.
[17] Eykholt et al., "Robust Physical-World Attacks on Deep Learning Visual Classification," 9-10.
[18] Akhtar and Mian, "Threat of Adversarial Attacks on Deep Learning in Computer Vision: A Survey," 26.
[19] Matt Turek, "Explainable Artificial Intelligence (XAI)," Defense Advanced Research Projects Agency. https://www.darpa.mil/program/explainable-artificial-intelligence

[20] Akhtar and Mian, "Threat of Adversarial Attacks on Deep Learning in Computer Vision: A Survey," 17.
[21] Wei Emma Zhang et al., "Adversarial Attacks on Deep Learning Models in Natural Language Processing: A Survey" (The University of Adelaide: Australia, Oct 2019), 18.
[22] Zhang et Al, "Adversarial Attacks on Deep Learning Models in Natural Language Processing: A Survey" 6.
[23] Ibid., 19.
[24] Nguyen et a Thanh Thi Nguyen et al., "Deep Learning for Deepfakes Creation and Detection" (Deakin University: Victoria, Australia, Sep 2019), 1.
[25] Jamieson, "Next Generation ISR Dominance Flight Plan Summary," 2.
[26] AFWIC, "Joint All-Domain Command and Control Sensing Grid," 7.
[27]Giovanni Sartor and Andrea Omicini, "The Autonomy of Technological Systems and Responsibilities for Their Use," in *Autonomous Weapons Systems: Law, Ethics, Policy,* ed. Nehal Bhuta et al. (Cambridge, UK: Cambridge University Press, 2016), 66.
[28] Ibid., 61.
[29] Ibid, 66.
[30] Matthew Price, Stephen Walker and Will Wiley, "The Machine Beneath: Implications of Artificial Intelligence in Strategic Decision-making," *PRISM*, Vol. 7, No. 4 (2018), 96-99.
[31] Hin-Yan Liu, "Refining Responsibility: Differentiating Two Types of Responsibility Issues Raised by Autonomous Weapons Systems," in *Autonomous Weapons Systems: Law, Ethics, Policy*, ed. Nehal Bhuta et al. (Cambridge, UK: Cambridge University Press, 2016), 341.
[32] Ibid., 344.
[33] Merel A. C. Ekelhof, "Lifting the Fog of Targeting," *Naval War College Review*, Vol. 71, No. 3 (Summer 2018), 63.
[34] Ibid., 81.
[35] Department of Defense, "Summary of the 2018 Department of Defense Artificial Intelligence Strategy," 8.
[36] Ibid., 15.
[37] Defense Innovation Board, "AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense" (31 Oct 2019), 12.
[38] David Danks (Department of Philosophy, Carnegie Mellon University), interview by the author, 10 Mar 2020.
[39] Ibid.
[40] Ibid.
[41] Ibid.
[42] Christian LaBiere (Department of Psychology, Carnegie Mellon University), interview by the author, 9 Mar 2020.
[43] Summary of "Future Directions in Human Machine Teaming Workshop," Office of the Under Secretary of Defense for Research & Engineering, Arlington, VA, 16-17 Jul 2019, 17.
[44] LaBiere, interview, 9 Mar 2020.
[45] Simon Marks and Abdi Latif Dahir, "Ethiopian Report on 737 Max Crash Blames Boeing," 9 Mar 2020. https://www.nytimes.com/2020/03/09/world/africa/ethiopia-crash-boeing.html
[46] Ibid.
[47] Summary of "Future Directions in Human Machine Teaming Workshop," 16-17 Jul 2019, 1.
[48] Leslie Blaha (Senior Research Psychologist, AFRL) interview by the author, 9 Mar 2020.
[49] Ibid.

[50] Daniel Keim et al, "Visual Analytics: Definition, Process, Challenges," in *Information Visualization*, ed. A. Keren et al. (Springer-Verlag, Berlin: 2008) 155-157.
[51] Blaha, interview, 9 Mar 2020.
[52] Keim et al, "Visual Analytics," 158.
[53] Alex Endert et al., "Semantic Interaction: Coupling Cognition and Computation through Usable Interactive Analytics," *Visualization Viewpoints* (Jul/Aug 2015), 6.
[54] Ibid., 7.
[55] Blaha, interview, 9 Mar 2020.
[56] Ibid.
[57] Peter Lamonica, "AFRL/RI Processing and Exploitation CTC," (brief, USAF Scientific Advisory Board, 6 Nov 2019), 5.
[58] Ibid., 12.
[59] Ibid. Blaha, interview, 9 Mar 2020.
[60] LaMonica, "AFRL/RI Processing and Exploitation CTC," 12.
[61] Ryan Luley (High Performance Systems, AFRL/RI), interview by the author, 21 Jan 2020.
[62] Ibid.
[63] Ibid.
[64] Danks, interview, 10 Mar 2020.
[65] Ibid.
[66] Ibid.
[67] Ibid.
[68] Ibid.
[69] Ibid.
[70] Blaha, interview, 9 Mar 2020. LeBiere, interview, 9 Mar 2020.
[71] Danks, interview, 10 Mar 2020.
[72] Defense Innovation Board, "AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense," 9.
[73] Danks, interview, 10 Mar 2020.
[74] Joe Pappalardo, "The Air Force Will Treat Computer Coding Like a Foreign Language," *Popular Mechanics* (13 Sep 2018). https://www.popularmechanics.com/technology/security/a23116594/air-force-coding-programming-language-mike-kanaan/
[75] Blaha, interview, 9 Mar 2020.
[76] Ekelhof, "Lifting the Fog of Targeting," 81.