NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

SUPPLEMENT TO CAPSTONE REPORT

TIP OF THE SPEAR: CAN SPECIAL FORCES LEAD THE WAY FOR MILITARY APPLICATIONS OF AI?

by

Joakim Riberg and Lars J. Selin

June 2021

Thesis Advisor:	Leo J. Blanken
Second Reader:	Vinnie Monaco

Approved for public release. Distribution is unlimited.



TIP OF THE SPEAR:

CAN SPECIAL FORCES LEAD THE WAY FOR MILITARY APPLICATIONS OF AI?

J. RIBERG & J. SELIN

Naval Postgraduate School

June 2021.



GETTING AHEAD IN THE AI RACE

As Artificial Intelligence (AI) continues to evolve with ever-increasing speed, especially in the commercial sector, the military community runs the risk of being left behind. This is especially problematic for the Special Operations Forces (SOF) worldwide as they often are the first to encounter new methods and technology employed by irregular adversaries. By studying contemporary research and interviewing current experts within the field of AI, this capstone report suggests which factors the SOF community must consider when acquiring and evaluating new AI capabilities. Avoiding the traps associated with AI can only be done in conjunction with other organizations, focusing on the right products and managing user expectation. As a solution, the report on which this presentation is based proposes a tentative model focused on risk mitigation and rapid procurement chains to allow for SOF to maintain its edge on the current battlefield.

WHAT IS THE PROBLEM WITH AI?

The current discussion about AI is often centered around how this new technology is plagued with accountability problems and biases. Faced with the idea that new technology might harm rather than help organizations, as soon as AI is mentioned, potential buyers often regard AI with skepticism and fear. This apprehension about new technologies often results in long investigations and assessments before a product can be implemented. However, there is little new about AI: depending on how AI is defined, it has been around for nearly 70 years (about as long as commercial flight).[1] Furthermore, the question about biases and accountability is deeply connected to the tasks the specific AI is expected to perform, rather than to the AI model itself. So how can we mitigate these potential pitfalls to shorten the time from idea to implementation?

Trust

For new technology to get accepted in society, humans must trust that the product performs as projected. However, in order to develop this trust, we must use, test, and try products. It does not matter if the specifications guarantee 100% effectiveness at a task if the product breaks down once every hour. If the product does not work, we will not trust nor use it. But, if a product actually performs its tasks and is reliable, we will gradually trust it, thereby allowing us to expect similar or better results from similar products. So how do we begin to build this trust?

➢ Expectation Management

Research has proven that our expectations reflect how much we trust products.[2] If the general feeling is that AI will revolutionize every aspect of our work tomorrow, many will feel that AI has let them down. In the same way that we don't expect a new rifle will make every round hit the target, we can't expect AI to perform miracles. Innovations and development seldom make giant leaps; instead, we should focus on constantly making small improvements.

➢ Trustworthy Partners

Although several organizations and companies promise exactly these quantum leaps in development, there is reason to be skeptical: a one-off sportscar made by an upstart company might have the better specifications, but you will probably buy a car from a more established brand, since you know that their cars usually are of high quality and comes with guarantees and established service sites. Hence, working with and through trustworthy partners within the commercial sector and academia, the military can maintain and create trust in the AI products it considers procuring.

➢ Education

Paramount in building the trust in AI is educating users. If users know how something works, they have a better understanding of what to expect. However, this does not mean that everyone must be an expert to use an AI-driven product; you do not have to be a certified mechanic to drive a car, but you do need a driver's license. Although there must be some experts skilled in repairing and analyzing the products, the military should tap the expertise that already exists in academia and the commercial sector to create innovation networks and benefit fully from the evolution of AI.

---

[1] John Haugeland, *Artificial Intelligence: The Very Idea*, 1st paperback ed. (Cambridge, MA: MIT Press, 1989), 176, https://ieeexplore-ieee-org.libproxy.nps.edu/xpl/ebooks/bookPdfWithBanner.jsp?fileName=6302870.pdf&bkn=6276821&pdfType=chapter.

[2] René F. Kizilcec, "How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, CA: ACM, 2016), 2393. https://doi.org/10.1145/2858036.2858402.

Managing the Risks of AI

Although AI promises new solutions for military applications, it also introduces several potential risks: Lack of *transparency* can mean that decisions are made on faulty assumptions, as the operator does not fully understand how the AI model reached a certain conclusion. The *security* of the AI-model itself must be maintained, as it can be attacked and corrupted through adversarial attacks. Similarly, the *data* used can be corrupted or stolen.



Figure 1: Visualization of the identified main components of AI security identified in the capstone report.

➢ Security

Apart from regular cyber security risks, two main threats to AI stand out: poisoning attacks and evasion attacks. Poisoning attacks target the ML algorithms and learning data (the ML model) by introducing new associations into the model to corrupt the model.[3] Evasion attacks are instead aimed at misguiding an otherwise working model into faulty classification by presenting an adversarial example.[4] With the possibilities offered through poisoning and input attacks of tricking ML models into missing or failing to detect threats, the consequences can be catastrophic, especially in a military context.

➢ Transparency

Transparency in AI comes from understanding; as Stoica et al. explain, "for a model to be fully understood, a human should be able to take the input data together with the parameters of the model and in reasonable time step through every calculation required to produce a prediction."[5] Simply put, a human has to be able to understand how and why an AI system functions. To benefit fully from transparency, organizations must focus efforts on educating operators and ensuring support by a core of highly skilled personnel capable of evaluating the potential benefits of incorporation of AI.

➢ Data

AI is driven by data. Without access to data to learn from or process, AI cannot function. However, since data constitutes the foundation for working AI, the integrity, safe storage, and access to data used become essential. Three main issues relate to the data used by AI: *biases*, *access*, and *cost*. Biases can occur from using bad data to train the AI or skew the analysis.[6] Access refers to how to obtain the needed data (own resources, commercial companies, or other government organizations.) Lastly, obtaining and maintaining data implies infrastructure costs to organizations.[7]

---

[3] Ion Stoica et al., *A Berkeley View of Systems Challenges for AI*, UCB/EECS-2017-159 (Berkley, CA: Electrical Engineering and Computer Sciences University of California at Berkeley, 2017), 5, https://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-159.pdf.

[4] Ilja Moisejevs, "Evasion Attacks on Machine Learning (or "Adversarial Examples")," Towards Data Science, July 2019, Evading Machine Learning, https://towardsdatascience.com/evasion-attacks-on-machine-learning-or-adversarial-examples-12f2283e06a1.

[5] Zachary C. Lipton, "In Machine Learning, the Concept of Interpretability Is Both Important and Slippery." *Queue* 16, no. 3 (May 2018),13. https://dl.acm.org/doi/pdf/10.1145/3236386.3241340.

[6] James Zou and Londa Schiebinger, "AI Can Be Sexist and Racist — It's Time to Make It Fair," *Nature* 559, no. 7714, July 2018, 324–326, https://doi.org/10.1038/d41586-018-05707-8.

[7] D. Sculley, Gary Holt, D. Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo and Dan Dennison., "Hidden Technical Debt in Machine Learning Systems," *NIPS* (2015), 1, https://www.semanticscholar.org/paper/Hidden-Technical-Debt-in-Machine-Learning-Systems-Sculley-Holt/1eb131a34fbb508a9dd8b646950c65901d6f1a5b

SPEEDING UP IMPLEMENTATION OF AI

Keeping up in the AI-arms race requires procurement policies and supply chains that maintain the relevance of the current situation. Fortunately, large parts of the SOF community around the world have rapid procurement chains, protocols for testing, and procedures in place for evaluation. However, regardless of how fast the procurement chain might be, the military must focus on the right products to avoid the risk of simple solutions getting bogged down in slow multi-organization procurement.

➢ Product and Capability Focus

Examining all possible ways to solve the problem at hand, instead of requiring the solution to incorporate AI, is a great first step to approach the problem. However, if the solution turns out to be based on an AI-driven product, it is important to focus on the actual tasks the product is expected to perform, instead of searching for potential pitfalls in the AI. Few have any objections about using a phone to support their everyday life, while many might object to using voice recognition AI to identify specific persons.

➢ Realistic AI Products

By avoiding products that promise autonomous actions and decision making outside of human control or intervention, it is possible to reduce the risk of getting caught up in ethical debates not yet been solved (and that might not be solved for a very long time), such as the question of self-driving cars. Instead, focusing on solutions that speed up decision making and actions while maintaining a human in the loop greatly enhances the possibilities of actual implementation.

➢ Front-end AI

A focus on user-friendly products of a “Plug-n-Play” character that are less reliant on r support of back-end systems to operate potentially limits the need for lengthy cooperative evaluations before implementation of the product. If the safety, maintenance, and risk of implementing the product can be accommodated by the organization, the potential time for implementation can be drastically reduced.



Figure 2: Simplified division of front- and back-end AI.

Evaluate the Current Organization

Assessing organizational capability is critical to make sure that the proposed solution will render the desired effects. When conducting the assessment, it is essential to address the question of capability:; Does the organization have the necessary knowledge to run and maintain an AI functionality? If not, then what is needed to acquire these critical capabilities?

Speed up the Process

When a *problem* or an *AI-solution* is identified or presented, the requirements of the actual *application* (not only focusing on the AI) should be identified. However, the depth in evaluation of the earlier identified parameters of transparency, security, and data might be greatly reduced if a proposed product is first assessed by its *application* and *risk*.

➢ Application

Where the product is intended to be employed physically and *How* the model will be used indicate the need for security, ethical considerations, training, infrastructure, and the required support and maintenance.

➢ Risk

Assessing the risk of failure or compromise of the AI product relates to *mission, tactical,* or *strategic* levels. If the failure will affect other parties or systems, their needs must be considered during the process.



Figure 3: Capstone model applied in a process.

If there is a need for a support structure outside the organization for the product to work efficiently, this might indicate that the solution is of a "*back-end*" character. If so, the need to work with the stakeholders of other connected systems and data must be considered.

If the solution is manageable within the organization, and more of a *"front-end"* character, the specific *risks* of the product should be considered. Higher risks would indicate that implementation might affect other parts of the military or organizations.

However, if there are no apparent issues identified, the SOF organization should be capable of implementing, testing, and evaluating the product within its own organization.

Why Is This Important?

Smaller organizations, like SOF, should focus on their strengths. In the case of SOF, this often entails the faster acquisition process rather than a high degree of technical expertise. Hence, collaborating with organizations capable of educating proficient AI professionals might be necessary for the more complicated applications. By determining the problem to be solved and applying the model, if the suggested solution contains an AI element, the organization should quickly realize if the capacity and capability to implement the product exists within the organization or if outside support is needed.

While larger and more complex solutions certainly will be necessary, the main point is to find a way to use, learn, and begin to trust AI solutions as quickly as possible. As more solutions containing AI are implemented, knowledge and capacity will grow within the organization, making it more receptive to new ideas in the future.

SUMMARY

Even though AI may continue to change several aspects of society, many questions still need addressing before AI gains the full trust of humanity. History provides many examples of innovations that altered the world; however, few innovations have revolutionized society within a few years. Most led to incremental changes to which the world has adapted. As the world continues to change, it seems likely that the integration of and advances in AI will continue to gradually affect the military as well. However, considering humans' reluctance to adapt to changes, SOF and their unique requirements might have an opportunity to remain at the tip of the spear. If a sound procurement process can be adopted, mitigating potential risks to and limitations of the organizations, it may be possible to increase knowledge of and trust in AI incrementally. Eventually, the broad base of acquired trust will allow for larger, more sophisticated solutions to be implemented with less resistance. As it turns out, learning to crawl might be the only way to win the AI race.

Goals

- Educate Users
- Build Trust in AI
- Manage Expectations

Actions

- Engage Trustworthy Partners
- Build Innovation Networks
- Review Slow Processes

References:

This document is intended as a visual and presentational supplement to the capstone report “Tip of The Spear: Can Special Forces Lead the Way in Military Applications of AI?” by J. Riberg & L. Johan Selin, June 2021, at the Naval Postgraduate school in Monterey, CA. It represents and uses, sometimes verbatim, the findings from the main capstone report. When applicable, other sources are referenced.

List of External References:

Haugeland, John. *Artificial Intelligence: The Very Idea*. 1st paperback ed. Bradford Books. Cambridge, Mass.: MIT Press, 1989. https://ieeexplore-ieee-org.libproxy.nps.edu/xpl/ebooks/bookPdfWithBanner.jsp?fileName=6302870.pdf &bkn=6276821&pdfType=chapter.

Kizilcec, René F. "How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface." In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2390–95. San Jose California USA: ACM, 2016. https://doi.org/10.1145/2858036.2858402.

Lipton, Zachary C. "In Machine Learning, the Concept of Interpretability Is Both Important and Slippery." *Queue* 16, no. 3 (May 2018):1–28. https://dl.acm.org/doi/pdf/10.1145/3236386.3241340.

Moisejevs, Ilja. "Evasion Attacks on Machine Learning (or "Adversarial Examples")," Towards Data Science, July 2019, Evading Machine Learning. https://towardsdatascience.com/evasion-attacks-on-machine-learning-or-adversarial-examples-12f2283e06a1.

Sculley, D, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, and Dan Dennison. Hidden Technical Debt in Machine Learning Systems, NIPS (2015). https://www.semanticscholar.org/paper/Hidden-Technical-Debt-in-Machine-Learning-Systems-Sculley-Holt/1eb131a34fbb508a9dd8b646950c65901d6f1a5b.

Stoica, Ion, Dawn Song, Raluca Ada Popa, David Patterson, Michael W. Mahoney, Randy Katz, Anthony D. Joseph, et al. *A Berkeley View of Systems Challenges for AI*. UCB/EECS-2017-159. Berkley, CA: Electrical Engineering and Computer Sciences University of California at Berkeley, 2017. https://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-159.pdf.

Zou, James, and Londa Schiebinger. "AI Can Be Sexist and Racist — It's Time to Make It Fair." Nature 559, no. 7714 (July 2018): 324–26. https://doi.org/10.1038/d41586-018-05707-8.