



Future Thinking: the Role Of Artificial Intelligence

November 17, 2017 *AFPC Defense Dossier Iss. 20, November 2017*

Related Categories: Cybersecurity and Cyberwarfare

The past several years have seen a remarkable transition of Artificial Intelligence (AI) from academia to practical use. This shift is beginning to transform every industry, is fundamentally changing many consumer services, and will have a profound impact on national security.

The current transformation is a half-century in the making. The 1956 Dartmouth Summer Research Project on Artificial Intelligence is properly credited with launching the field of AI.[1] However, it was the insightful 1960 publication by J.C.R. Licklider that outlined the prospect for computers to "facilitate formulative thinking as they now facilitate the solution of formulated problems, and to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs." [2] This symbiotic interaction between humans and computers is a foundational principle of the new ubiquity of AI.

In the decades after the Dartmouth study, the Department of Defense sponsored key work in many AI disciplines - among them speech recognition, natural language understanding and neural networks. These investments have borne profoundly transformational fruit in recent years, and promise to do so for years to come. In tandem, parallel investments by government, industry and academia in high performance computing, machine learning and more recently cloud services have provided the key technical foundation to integrate and apply AI to real-world systems. It is now our task to decide how to bring the next phase of this technology forward for the benefit of national security.

BRAVE NEW WORLD

Two seminal events of recent years marked the beginning of the AI-race to commercialization that is underway today.

In 2011, the IBM open-domain question-answering system known as Watson beat the two highest ranked quiz show players in a nationally televised two-match Jeopardy! contest.[3] This was the first live demonstration of integrated computational linguistics, information retrieval, knowledge representation and reasoning, and machine learning in an unstructured environment: questions and answers in the language of Jeopardy!

The following year, the field of Deep Learning was launched with the publication of a seminal paper from the University of Toronto that used convolutional nets to demonstrate almost half the error rate for object recognition, and precipitated the rapid adoption of deep learning by the computer vision community.[4]

These developments, coupled with the availability of massive data, embedded analytics and wide availability of GPUs that make parallel processing ever faster, cheaper, and more powerful, are today propelling AI into almost every industry. AI is now being used to derive insight from massive data sets that are dynamic, cluttered and in many cases ambiguous. And while these trends are, to a large extent, confined to the commercial and scientific space, the ability of AI to identify emerging trends and forecast potential courses of action is bound to have a similarly profound impact on national security.

THE IMPORTANCE OF INDUSTRIAL AI

McKinsey & Co. recently reported[5] that between \$26 and \$39 billion was invested in AI by companies in 2016 alone. Sixty-six percent of that investment was made in the United States. Industrial AI, in other words, is undoubtedly a growth industry; venture funding in this space alone has been growing at an average of 40 percent annually for the past 3 years; and 61 percent of the 3,000 companies interviewed by the consultancy group said that they are either AI adopters or partial adopters. This rapid growth in investment and adoption, in turn, has been driven by a number of recent machine learning success stories, such as the estimates of online movie site Netflix[6] that it now saves some \$1 billion annually thanks to its AI-based video recommendation system.

Industrial AI solutions vary greatly because of the differing needs of customers. But they all share several features: the ability to ingest vast amounts of data from varied sources; the ability to curate that data for reliability, provenance, value, shelf-life, etc.; the ability to extract meaningful, relevant information; and the ability to convey this information to a user for decision making or to another device for action. For example, IBM's Watson Health system[7] integrates these features to provide improved healthcare outcomes by analyzing volumes of data that no human could possibly read, and digesting and prioritizing situationally relevant information to healthcare professions which can then use it to make treatment decisions.

For the Pentagon, in turn, leveraging this market momentum should be a top priority.

TOWARD A PREEMPTIVE NATIONAL SECURITY INTELLIGENCE ENVIRONMENT

A new approach is likewise needed to address increasingly complex nation-state threats and an emerging era of decentralized dangers that have blended into society and are evolving faster than our current ability to collect, model, understand and interdict them. Responding to these threats has pushed the U.S. into a more reactive posture - one in which we are responding to events after they happen. This new landscape poses great danger to U.S. national security and American interests abroad.

AI technology offers the American military an opportunity to enable the next era of national security by building a preemptive national security intelligence environment that is centered on forecasting, shaping and disrupting national security threats well before they fully present themselves.[8] These advanced services will give our nation the ability to detect the emergence of threats by modeling the processes that build adversarial capability "maturity" over time and by tracking their progress even where the telltale signs may be only partially visible or deliberately obfuscated. This approach could help inform the following new capabilities.

1) *Persistent Over-Watch (moving analytics from offline to in-line)*: Automated data analytics, machine learning, and composable cognitive environments that constantly monitor and extract real-time insight from live data to detect emerging scenarios as they are beginning to form, project how those scenarios could evolve and recommend the most effective preemptive near-term courses of action.

2) *Simulation-Based Forecasting (preemptively shaping the threat)*: A set of simulation and game theory services with integrated military, political, economic and social models that forecast how to best shape an adversary's calculus and recommend preemptive long-term courses of action to shape the outcome.

3) *Immersive Data Environments (visualizing the threat)*: Immersive environments that allow analysts and leaders to visualize, interact with and understand complex intelligence data, evolving scenarios, live metrics and evolving courses of action. This environment will enable analysts and operators to collaborate and evaluate "what if" scenarios and shape the environment by observing new data relationships.

4) *Accelerated DevOps (Innovation, speed & agility)*: A development environment that allows the DoD and the intelligence community (IC) to develop, onboard, provision and utilize new capabilities faster than the adversary with a scaled network effect. This composable environment will be key to protecting the technological gains made by the DoD.

5) *Distributed Decision Making*: By embedding machine learning decision-support systems alongside sensors, future intelligence, surveillance and reconnaissance (ISR) systems will support distributed decision-making for theater commanders requiring more rapid responses to threat signatures.

CHALLENGES AHEAD

Although we have recently seen dramatic improvement in AI capabilities in recent times, as well as a rapid growth in the adoption of applications that use them, exploiting AI to its full potential still faces numerous difficult challenges.

Today, AI can, with high reliability, identify dogs and cats in images, and it is beginning to be able to identify multiple general objects in an image. But it is still very far from converting a picture into a story, or predicting from an image of a falling coffee mug that something is about to break. These kinds of tasks, known as Artificial General Intelligence[9], involve being able to understand and reason about the real world, and are still stubbornly beyond our grasp. Similarly, for natural language understanding, AI can reliably convert between speech and text, and is beginning to reliably translate between languages. But deeply understanding and reasoning about, for example, the nuances of a legal document is still an unsolved problem.

Trust is another major challenge for AI systems, whose learning algorithms typically do not yield interpretable models and therefore make it very difficult to verify that they always do what is expected of them. For example, it is impossible to test every possible driving situation for a self-driving car, so how does one know the system will respond appropriately to every new situation? To develop trust, practitioners rely on statistical evaluations of AI systems with millions of test cases, or more. However, the complexity of these systems and the sheer volume of data can thwart the best of intentions. Consider the case of the Google Photos application which erroneously classified images of people as gorillas and which led to serious public relations problems for Google.[10] In military applications, the mistakes of an AI system could have extremely dire outcomes. Satisfactory methods for ensuring AI system reliability in mission critical situations do not exist today. Related to this is the recently identified weakness for "adversarial examples" exhibited by Deep Learning.[11]

Methods exist to easily create data that makes deep learning systems do the wrong things, even without knowledge of the details of the system in question. In one specific example, it is possible to modify a picture of a bus such that it gets recognized as an ostrich, even though the modification is so slight that a human cannot visually detect the change, and certainly cannot see an ostrich in the image. Worryingly, adversarial examples are not limited to buses and ostriches; they can be demonstrated much more generally. This known weakness leaves AI systems open to an array of attacks by malicious actors. Countering this vulnerability is an active field of research.

THE WAY FORWARD

Even with its faults, AI technology holds tremendous promise for extending U.S. military effectiveness, for acting as a force multiplier, and for providing a third offset mechanism that will move us to a more proactive posture, allowing us to shape threats before they fully manifest. This capability is especially important as the nature of the military's challenges continues to evolve, and we continually strive to improve in the most cost-effective manner possible. It is imperative that we take full advantage of these opportunities.

Moving forward on this path requires us to provide sufficient investment to advance the core AI technologies required, many of which are still active areas of research. We also must carefully think through how these technologies will be integrated with existing systems and operational methods so they will provide maximal benefit and not detract from what already works well. Finally, we must develop rigorous methods for validating the effectiveness, reliability and vulnerabilities of AI systems in order to prevent unintended consequences.

Despite the significant challenges that lie ahead, AI's potential to change the military landscape presents the U.S. with a unique opportunity. We ought to take advantage while the time is right.

Mr. Zachary Lemnios leads Physical Sciences and Government Programs, globally across IBM Research, to extend fundamental scientific understanding and breakthroughs that enable the future of information technology. Prior to joining IBM, Mr. Lemnios was confirmed as Assistant Secretary of Defense (Research & Engineering) by the United States Senate. He also served as the Chief Technology Officer of MIT Lincoln Laboratory and led the development of advanced technologies in support of national security.

Dr. Michael Perrone is the AI Partnership Program Director at IBM Research. Michael was instrumental in the establishment of IBM's Cognitive Horizons university research network, and is now actively pursuing opportunities to enhance military effectiveness with IBM AI technologies, and to extent machine learning to massive-scale computing. His previous work resulted in multiple international awards, and an IBM Master Inventor award.

NOTES:

- [1] John McCarthy, Marvin L. Minsky, Nathaniel Rochester, Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," *AI Magazine*, August 31, 1955, <https://www.aaai.org/ojs/index.php/aimagazine/article/view/1904/1802>
- [2] J. C. R. Licklider, "Man-Computer Symbiosis," *IRE Transactions on Human Factors in Electronics*, March 1960, <http://worrydream.com/refs/Licklider%20-%20ManComputer%20Symbiosis.pdf>
- [3] E.W. Brown et al., "This is Watson," *IBM Journal of Research and Development* 56, iss. 3.4, Summer 2012, 1-17., <http://ieeexplore.ieee.org/xpl/tocresult.jsp?isnumber=6177717>
- [4] Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *NIPS '12 Proceedings of the 25th International Conference on Neural Information Processing Systems*, December 3, 2012, <https://papers.nips.cc/paper/4824-imagenetclassification-with-deep-convolutional-neural-networks>
- [5] Jacques Bughin, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlstrom, Nicolaus Henke, Monica Trench, "How Artificial Intelligence Can Deliver Real Value to Companies," McKinsey Global Institute, June 2017, <http://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/how-artificial-intelligence-can-deliver-real-value-to-companies>
- [6] Carlos A. Gomez-Urbe, Neil Hunt, "The Netflix Recommender System: Algorithms, Business Value, and Innovation," *ACM Transactions on Management Information Systems (TMIS)* 6, iss. 4, January 2016, 1-19, <http://dl.acm.org/citation.cfm?id=2843948>
- [7] IBM, "Empowering Heroes, Transforming Health," IBM.com, n.d., <https://www.ibm.com/watson/health>
- [8] N. Halim, Z. Lemnios, G. Loftin, "Staying Ahead of the Threat: A preemptive National Security Intelligence Environment".
- [9] "Artificial General Intelligence," Wikipedia.org, n.d., https://en.wikipedia.org/wiki/Artificial_general_intelligence
- [10] Jessica Guynn, "Google Photos Labeled Black People 'Gorillas'," *USA Today*, July 1, 2015, <https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photosidentify-black-people-as-gorillas/29567465>
- [11] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, Rob Fergus, "Intriguing Properties of Neural Networks," Cornell University: Computer Vision and Pattern Recognition, February 19, 2014, <https://arxiv.org/pdf/1312.6199.pdf>