




HDIAAC

Homeland Defense & Security
Information Analysis Center



Artificial Intelligence

State of the Art Report

DISTRIBUTION A. APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

**STATE OF THE ART REPORT:
Artificial Intelligence and Machine Learning for Defense Applications**

Gregory Nichols, Homeland Defense & Information Analysis Center
Sue Ellen Haupt, National Center for Atmospheric Research
David John Gagne, National Center for Atmospheric Research
Anthony A. Rucci, GRIDSMART Technologies, Inc.
Joel Hewett, Homeland Defense & Security Information Analysis Center
Gopikrishna Deshpande, Auburn University
Pradyumna Lanka, Auburn University
Susan A. Youngblood, Auburn University

2017

Contract Number: FA8075-13-D-0001

Sponsorship Statement:

HDIAC is a DoD sponsored Information Analysis Center (IAC) with policy oversight provided by the Under Secretary of Defense for Research and Engineering (USD (R&E)), and administratively managed by the Defense Technical Information Center (DTIC). HDIAC is operated by Information International Associates (IIa) in Oak Ridge, Tenn. Reference herein to any specific commercial products, processes or services by trade name, trademark, manufacturer or otherwise, does not necessarily constitute or imply its endorsement, recommendation or favoring by the United States government or HDIAC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or HDIAC and shall not be used for advertising or product endorsement purposes.

About This Publication:

Copyright 2018 by IIa. This State of the Art Report was developed by IIa under HDIAC contract FA8075-13-D-0001. The government has unlimited free use of and access to this publication and its contents in both print and electronic versions. Subject to the rights of the government, this document (print and electronic versions) and the contents contained within it are protected by U.S. copyright law and may not be copied, automated, resold or redistributed to multiple users without the written permission of HDIAC.

Distribution A:

Approved for Public Release; Distribution Unlimited

This page intentionally left blank.

Contents

AUTHORS	4
ACRONYMS	6
EXECUTIVE SUMMARY	7
1 INTRODUCTION	10
2 USE OF ARTIFICIAL INTELLIGENCE IN ENABLING RENEWABLE ENERGY	23
Introduction, 23	
Forecasting Renewable Energy and AI, 25	
A Look to the Future, 34	
Summary and Conclusions, 36	
3 BIOMETRICS, CRITICAL INFRASTRUCTURE PROTECTION, AND HOMELAND DEFENSE AND SECURITY	45
4 CHEMICAL, BIOLOGICAL, RADIOLOGICAL, AND NUCLEAR DEFENSE/WEAPONS OF MASS DESTRUCTION	52
5 ARTIFICIAL INTELLIGENCE APPLICATIONS TO MILITARY MEDICINE	59
Introduction, 59	
Applications of AI in Mental Health Disorders and Brain Injuries, 60	
Data Sharing and Big Data Neuroimaging Initiatives, 63	
Conclusion, 64	
6 CONCLUSION	71

Authors

Gregory Nichols, MPH, CPH

Homeland Defense & Information Analysis Center

Gregory Nichols is a subject matter expert for the Homeland Defense & Security Information Analysis Center (MPH, University of Tennessee). Previously, he managed the Nanotechnology Studies Program at ORAU in Oak Ridge, Tennessee, where he provided expertise on nanotechnology-related topics and conducted research. Prior to ORAU, Nichols spent 10 years in various health care roles, including five years as a Hospital Corpsman in the U.S. Navy. He has published and presented on a variety of topics including nanotechnology, public health, and risk assessment.

Sue Ellen Haupt, Ph.D.

National Center for Atmospheric Research

Sue Ellen Haupt is senior scientist and director of the Weather Systems and Assessment Program of the Research Applications Laboratory of National Center for Atmospheric Research. She is also director of education of the World Energy and Meteorology Council and a councilor of the American Meteorological Society. She previously headed the Department of Atmospheric and Oceanic Physics at the Applied Research Laboratory of The Pennsylvania State University where she remains an adjunct professor. Haupt is an expert in renewable energy, boundary layer meteorology, large-scale atmospheric dynamics, numerical methods, artificial intelligence methods, and computational fluid dynamics.

David John Gagne, Ph.D.

National Center for Atmospheric Research

David John Gagne is an Advanced Study Program Postdoctoral Fellow at the National Center for Atmospheric Research in Boulder, Colorado (Ph.D., University of Oklahoma). His research focuses on developing and evaluating machine learning systems for high-impact weather forecasting problems. Gagne has coauthored 10 journal articles in both meteorology and machine learning. He developed a gridded forecasting system for solar energy, administered the American Meteorological Society Solar Energy Prediction Contest, and integrated machine learning models into wind and solar forecast systems.

Antonio A. Rucci

Information International Associates

Antonio A. Rucci is a retired Counterintelligence Special Agent and U.S. Army Warrant Officer with more than 33 years of technical security experience. He is currently the director of information security and threat intelligence at GRIDSMART Technologies, Inc. Rucci ran technical espionage and computer crimes investigations for a good portion of his operational career and served as the Counterintelligence Operations Officer at The White House, under two

administrations. Rucci continues to serve national security efforts in his current capacity with Ila, building a robust threat intelligence and information security vertical, supporting both private sector and government interests.

Joel Hewett

Homeland Defense & Security Information Analysis Center

Joel Hewett is a subject matter expert with the Homeland Defense & Security Information Analysis Center (M.S., Georgia Institute of Technology). His book-length study of the offshore oil and gas program in the U.S., *The Secret of the Sea*, will be published by the Department of the Interior in 2018. Hewett was the Georgia Institute of Technology inaugural Melvin Kranzberg Fellow. He also holds an A.B. in literature from Davidson College, where he was a John Montgomery Belk Scholar.

Gopikrishna Deshpande, Ph.D.

Auburn University

Gopikrishna Deshpande is an associate professor at the Magnetic Resonance Imaging Research Center within the Department of Electrical and Computer Engineering at Auburn University (Ph.D., Georgia Institute of Technology; M.S., Indian Institute of Science Bangalore, India). He is the recipient of Auburn University's 2016-2017 President's Outstanding Collaborative Unit Award, has published more than 60 peer-reviewed articles on neuroimaging, signal/image processing, and machine-learning, and has been awarded approximately \$4 million in federal research grants. His work on canine-imaging and post-traumatic stress disorder have been prominently featured in national and international media.

Pradyumna Lanka

University of California, Merced

Pradyumna Lanka is pursuing a Ph.D. in psychological sciences at the University of California, Merced (M.S., Auburn University; B.S., Birla Institute of Technology and Science, Pilani, India). His research interests include functional neuroimaging, particularly functional magnetic resonance imaging and functional near-infrared spectroscopy; signal/image processing; pattern recognition; and machine learning. Lanka is especially interested in developing tools for automated disease diagnosis.

Susan A. Youngblood, Ph.D.

Auburn University

Susan A. Youngblood is an associate professor of technical and professional communication in the Department of English at Auburn University (Ph.D., Texas Tech University). Youngblood researches online communication choices, particularly in situations wherein sharing information poses safety or security risks (e.g., risks of domestic terrorism), creates barriers to accessibility for vulnerable populations, or is complicated by perceptions.

Acronyms

AC	Artificial Conscience
ADHD	Attention-deficit/hyperactivity Disorder
AnEn	Analog Ensemble
AFRL	Air Force Research Laboratory
AI	Artificial Intelligence
ALPAC	Automatic Language Processing Advisory Committee
ARIMA	AutoRegressive Integrated Moving Average
ARPA	Advanced Research Projects Agency
ASD	Autism Spectrum Disorder
AVATAR	Automated Virtual Agent for Truth Assessments in Real-Time
CBRN	Chemical, Biological, Radiological, and Nuclear
CIP	Critical Infrastructure Protection
CPU	Central Processing Unit
CV	Computer Vision
DARPA	Defense Advanced Research Projects Agency
DART	The Dynamic Analysis and Replanning Tool
DHS	Department of Homeland Security
DoD	Department of Defense
DPV	Distributed Photovoltaic
ENSO	El Nino/Southern Oscillation
FEDOR	Final Experimental Demonstration Object Research
GCM	Global Climate Model
GW	Gigawatt
HDIAC	Homeland Defense & Information Analysis Center
HDS	Homeland Defense and Security
IARPA	Intelligence Advanced Research Projects Activity
ISO	Independent Systems Operator
LRASM	Long Range Anti-Ship Missile
MOS	Model Output Statistics
MOU	Memorandum of Understanding
mTBI	Mild Traumatic Brain Injury
NCAR	National Center for Atmospheric Research
NRC	National Research Council
NWP	Numerical Weather Prediction
ONR	Office of Naval Research
PCS	Post-Concussion Syndrome
pdf	Probability Density Function
PTSD	Post-Traumatic Stress Disorder
RCM	Regional Climate Model
RPD	Recognition-primed Decision
RSFC	Resting-state Functional Connectivity
SAM	Surface-to-Air Missile
SOM	Self Organizing Map
TPU	Tensor Processing Unit
UNICRI	United Nations Interregional Crime and Justice Research Institute
USPACOM	United States Pacific Command
WMD	Weapons of Mass Destruction
WRF	Weather Research and Forecasting

Executive Summary

The Homeland Defense & Security Information Analysis Center (HDIAC) develops State of the Art Reports on relevant, key topics for the Department of Defense (DoD). One of the most significant topics within the Defense community is artificial intelligence (AI). AI and some of its branches, primarily machine learning, were prominent in the 2014 DoD Third Offset Strategy, which guided DoD's approach to combat operations and defensive posture [1]. AI research has also become pervasive in academia and nearly all industrial sectors. Some of this research produces outcomes being used by DoD and also highlights new risks and challenges that could have implications for the DoD regarding how it conducts operations in the future.

AI research in general, as well as its use in military applications, has its roots in the 1950s. Alan Turing, a British mathematician and computer scientist, asked the question, "Can machines think [2]?" This introduced a realistic tenor to the field of AI, a subject once isolated to the realm of science fiction. Almost immediately after the publication of Turing's article, the Office of Naval Research (ONR) funded early work into artificial neural networks, which form the basis of machine learning and are required for deep learning [3]. Subsequently, AI efforts, both in defense and civilian research, were led by the Advanced Research Projects Agency (ARPA, which was later renamed Defense Advanced Research Projects Agency [DARPA]), which continues to this day [4]. However, additional research is now being conducted in the private sector, which poses its own set of challenges. Just over half a century later, the concept of machine thought is rapidly approaching reality. Recent advances in the field demonstrate new AI capabilities, such as sight, speech, and even moral/ethical reasoning [5-7]. These advances all show a common theme that machines are able to learn, adapting or amending their predetermined programming based on new input. The ability for machines to learn can be incorporated into weaponry and other applications used within DoD to support DoD strategy. AI is redefining the way wars are fought and won, as has been the case with many emerging technologies throughout history.

Since the 1970s, the bulk of DoD efforts regarding AI have focused on applied research [4]. Accordingly, this report focuses on current and potential applications of AI—not basic research. In many ways, this narrower focus makes it easier to describe the proposed capabilities and current capacity of AI. Additionally, new research is published almost daily, making it challenging to stay up-to-date on essential algorithms, coding principles, chip processing speeds, and other pertinent information needed to develop AI.

This report focuses on applications of AI that are relevant to DoD and other governmental agencies that share similar goals, such as the U.S. Department of Homeland Security (DHS) and the Intelligence Community. HDIAC collaborated with subject matter experts to obtain interpretations on relevant applications of AI to DoD and other agencies that could similarly incorporate AI into their operations. Discussions narrowed on relevant applications that have been in use for the past three years as well as advances that could become commercialized within the next 18–24 months. Given the vast amount of information, this report is not all-inclusive and is only a survey of prominent developments.

Although AI is often thought of as a part of computer science and typically viewed from a cyber perspective, applications of AI are relevant to all eight HDIAC focus areas. This report covers applications to Alternative Energy; Biometrics; Chemical, Biological, Radiological, and Nuclear (CBRN) Defense; Critical Infrastructure Protection (CIP); Cultural Studies; Homeland Defense and Security (HDS); Medicine; and Weapons of Mass Destruction (WMD). In some cases, the discussions group together areas that either overlap well or where information on AI applications is thin, and it made more sense to integrate relevant areas in order to provide a more robust perspective.

References

1. Lange, K. (2016, March 30). 3rd Offset Strategy 101: What it is, what the tech focuses are. DoDLive. Retrieved from <http://www.dodlive.mil/2016/03/30/3rd-offset-strategy-101-what-it-is-what-the-tech-focuses-are/>
2. Turing, A. M. (1950). I.—Computing machinery and intelligence. *Mind, LIX*(236), 433-460. doi:10.1093/mind/lix.236.433
3. Miller, W., Mckenna, T., & Lau, C. (1992). Office of Naval Research contributions to neural networks and signal processing in oceanic engineering. *IEEE Journal of Oceanic Engineering, 17*(4), 299-307. doi:10.1109/48.180298
4. National Research Council. (1999). Funding a revolution: Government support for computing research. Washington, DC: The National Academies Press. doi:10.17226/6323
5. Metz, C. (2017, June 02). The unsettling performance that showed the world through AI's eyes. Retrieved from <https://www.wired.com/2017/04/unsettling-performance-showed-world-ais-eyes/>
6. DepMind. (n.d.). WaveNet: A generative model for raw audio. Retrieved from <https://deepmind.com/blog/wavenet-generative-model-raw-audio/>
7. Vicky Charisi, V., Dennis, L., Fisher, M., Lieck, R., Matthias, A., Slavkovik, M., Loh (Sombetzki), J., Winfield, F.T., & Yampolsky, R. (2017). Towards moral autonomous systems. Retrieved from <https://arxiv.org/abs/1703.04741>

1

Introduction

Gregory Nichols, Homeland Defense & Security Information Analysis Center

Background

One of the most significant topics in the defense community is AI. As illustrated by the establishment of a U.S. national AI research strategy last year [1], high expenditures on AI research within DoD [2], announcements by other nations (such as China and Canada) to be leaders in AI [3,4], and the incorporation of AI into weapons systems over the past five years, AI is an important topic of discussion within the DoD community. Although typically considered within the realms of computer science or cybersecurity, AI may also be considered in terms of all eight HDIAC focus areas (see Figure 1), making this a meaningful topic of discussion.



Figure 1. Examples of AI applications across HDIAC focus areas.

Overview

Intelligence refers to intellectual functioning [5]. It is the ability to learn or understand or to deal with new or trying situations [6], and also encompasses a capacity for logic, understanding, self-awareness, learning, emotional knowledge, planning, creativity, and problem-solving [7,8]. Intelligence is exhibited by a variety of animal species (most notably humans). The true test of intelligence is unclear, but, most likely, intelligence involves the creation and use of persistent memories (as opposed to computation that does not involve some aspect of learning),

excluding purely autonomic sense-reaction responses. Based on these defining characteristics, it is possible to create a machine capable of learning from previously recorded experiences—thus recreating intelligence. These principles are the foundation of the theory that machines could be programmed to learn, thereby becoming intelligent.

AI is the intelligence exhibited by machines. Goals related to AI include the development of intelligent agents that can recreate in machines those characteristics of intelligence attributed to humans, such as reasoning, planning, learning, natural language processing, emotional expression, and creativity [9,10].

The concept of AI involves principles from many disciplines, including philosophy, mathematics, engineering, and computer science. The notion of AI traces its roots to ancient Greece when the philosopher Aristotle first formalized the fundamentals of modern logic. Logic is a philosophical and mathematic discipline that is concerned with proper reasoning [11]. This process can be used to deconstruct the world into small, understandable components that can be used to build more complex analyses and evaluate connections within and between any system. Over the past 200 years, these principles of logic have converged with the principles of mathematics to more accurately reflect how humans process and comprehend the natural world, much of which has been responsible for the creation of programming and the foundation of computer science. Contributions from philosophers and mathematicians (such as Alonzo Church, Gottlob Frege, Kurt Gödel, Bertrand Russell, Ludwig Wittgenstein, and Alan Turing) led to the development of the platform needed to develop AI, primarily computability.

Because the goal of AI is to mimic the complex output of the human brain, a significant portion of AI research focuses on what it means to be human and how the essence of humanity can be simulated by a machine. In addition to logic, work in AI requires knowledge of the philosophy of mind (a branch of philosophy that deals with the nature of thought and consciousness) [12] as well as psychology (the study of the mind and behavior) [13]. Because of this, most research has focused on recreating the machinery of the brain in order to emulate natural pathways such as neurons that propagate electrical signals and facilitate thought and action. Artificial neural networks are a key component of AI development, particularly within the branch of AI known as machine learning. Machine learning involves pattern recognition, which is taught to a machine by allowing it to process thousands, if not millions, of images until it can recognize similarities and differences between objects in order to determine what is a match and what it is not. This process is innate to humans, illustrated by our ability to remember faces, identify people, and simultaneously recognize similarities while understanding differences. However, machines do not innately possess pattern recognition, so they must be taught.

Machines also struggle with learning because they cannot “remember” in the same manner as humans. Learning and understanding is derived from experience and remembrance of the past while building toward a more complex task. Machines suffer from what is known as catastrophic forgetting, which means they are capable of learning one task at a time but are not completely able to put these tasks together in a sequence on their own [14].

The aforementioned challenges limit machines to performing specific tasks as designed, an approach referred to as weak or narrow intelligence. All current applications of AI are in this area. The ultimate goal is to develop strong or general AI that more broadly reflects the complex range of tasks and emotions demonstrated by humans, eventually achieving sentience, or self-awareness.

Interest in and funding related to AI has waxed and waned over the years. At its inception in the 1950s, AI research began with expert systems designed by experienced professionals who programmed machines to do tasks based on specific knowledge. However, the unprecedented spending associated with this research halted in the 1970s due to the unpleasant revelations of the Lighthill report [15] and did not pick up again until the 1980s [16]. The development of better algorithms, chips with faster processing speeds, and increased private sector funding have spurred AI development since approximately 2006 [17], and researchers are now transiting to statistical-based programming. Funding slowed down once more until the recent bump in spending again, primarily spurred by private sector companies (such as Google, Apple, Baidu, and Microsoft), which spent an estimated \$26-\$39 billion on AI in 2016 [18], as well as a shift in DoD strategy [19].

History of AI in Defense, the Intelligence Community, and Homeland Security

Modern AI interest evolved from the work of Alan Turing who assisted in the development of a machine used to crack the German “Enigma” code during World War II [20]. Following the war, Turing worked on designs for a forerunner to the modern computer (the Automated Computing Engine) and published a paper in 1950 titled “Computing Machinery and Intelligence” in which he asked the question: “Can machines think [21]?” The paper explained how AI is possible and refuted many common arguments against the concept that existed at that time, establishing the Turing Test (see Figure 2) as the process of evaluating the ability of a machine to “exhibit intelligent behavior equal to or indistinguishable from a human [21].” In 1956, John McCarthy organized the *Dartmouth Summer Research Project on Artificial Intelligence* in which he first coined the term “artificial intelligence” to describe this machine-demonstrated reasoning [22].




Figure 2. A basic interpretation of the Turing Test, in which one player attempts to determine which of the other players is human and which is a computer. Adapted from Copeland [23].

Following this project, the pace of AI research quickened, interest grew, and funding rapidly increased. After Turing's success using computing to decipher the "Enigma" code, British and American militaries were interested in the use of more advanced computing methods for not only intelligence gathering but also for developing weapons systems. In 1957, the Office of Naval Research funded the development of one of the first artificial neural networks, the Mark 1 perceptron (see Figure 3). Early efforts focused on software development, but the Mark 1 was designed for image recognition [25]. It had an array of 400 photocells. Weights were encoded in potentiometers, and weight updates during learning were performed by electric motors [25].




Figure 3. Mark 1 perceptron [24]. (Released)

Support for AI began to wane in the 1960s due to a perceived lack of substantial developments in relation to the cost of research. This decrease in support was followed by cyclic periods of hype and then abandonment, which has since continued, referred to as AI winters (lack of support) and AI springs/summers (renewed support) [26]. For example, one of the primary research areas of AI had been in machine translation. The National Research Council (NRC) expressed concern regarding the lack of progress in the field of AI and, in 1964, formed the Automatic Language Processing Advisory Committee (ALPAC) to investigate. ALPAC concluded that "machine translation was more expensive, less accurate, and slower than human translation [27]." Following the release of the report, the NRC suspended support for machine translation, which continues to be a challenging research problem.

From 1964 to 1966, at the same time ALPAC was established, a breakthrough in natural language processing occurred at the Massachusetts Institute of Technology (MIT) Artificial Intelligence Laboratory. A computer program named ELIZA became the first program to arguably pass the Turing test by crudely mimicking a psychotherapist [28,29]. However, in 1973, the United Kingdom-

based Science Research Council commissioned Sir James Lighthill of Cambridge University to evaluate the state of AI research throughout the U.K. His conclusions were released in an infamous report, commonly referred to as the Lighthill Report, in which he determined that “in no part of the field have discoveries made so far produced the major impact that was then promised [30].” Upon its publication, AI research halted throughout the U.K., except for efforts at three universities, and the shockwaves rippled to the United States.

Throughout the 1960s and early 1970s, most AI research in the U.S. had been funded by ARPA and performed at Carnegie Mellon University, Stanford University, and MIT [16]—all of which continue to be leaders in AI. Fears of a strengthening military-industrial complex fueled by American involvement in Vietnam and a perception of uncontrolled and misguided defense research led Congress to pass the Mansfield Amendment in 1969 [16]. The Mansfield Amendment restricted DoD research to be “direct and apparent” to specific military functions and operations, significantly affecting DARPA-supported AI research, which was primarily basic research [16,31]. Later, the the Mansfield Amendment of 1973 affected DARPA funding even more [32]. Shortly thereafter, the Lighthill Report was released, prompting DARPA to create a panel of AI experts, the American Study Group, which reached conclusions similar to the Lighthill Report. Therefore, DARPA-funded AI research shifted to mission-direct research [16]. In 1974, the DARPA-funded Speech Understanding Research (SUR) Program at Carnegie Mellon University was cancelled. The SUR Program was intended to recognize voice commands given by a pilot; however, a demonstration revealed that while the program recognized words in English, they had to be spoken in a particular order to execute a command. DARPA determined the intended results were not achievable, and the program was dismantled [16]. However, AI research picked up once again a decade later in 1983 when DARPA created the Strategic Computing Initiative (SCI)—a decade-long initiative that yielded significant work in AI [16]. Both the work that stemmed from SCI and the shift from basic to applied research set the stage for the computing revolution in the U.S. and created a foundation for DoD’s Third Offset Strategy.

While the U.S. and U.K. governments struggled with AI research, breakthroughs in AI research continued. The Soviet Union developed the P-700 Granit (SS-N-19 Shipwreck) in the 1970s, which is one of the first cruise missiles with an AI tracking system [33]. And in 1977, MYCIN, an AI system, was used at Stanford University to treat blood infections. It was one of the first AI systems to be put in use for a specific function [34]. As the years progressed, previous work in AI at DARPA began to pay off. The Dynamic Analysis and Replanning Tool (DART) was first introduced in 1991 during Operation Desert Storm to optimize logistics. DART was so successful that within four years, it offset the cost of the previous 30 years of all DARPA AI research [35]. Furthermore, one of the key breakthroughs in AI occurred in 1997, when IBM’s Deep Blue defeated chess grandmaster Gary Kasparov, demonstrating that it is possible for machines to out-think people [36].

The beginning of the new millennium saw a resurgence in AI research, particularly among the defense and intelligence communities. Throughout the 2000s, DARPA funded several AI-oriented programs. From 2003 to 2008, the Cognitive Assistant that Learns and Organizes project, known as CALO, established the foundation for the personal assistant, Siri, used in all Apple devices [37]. In 2004, DARPA

launched the DARPA Grand Challenge, a prize contest for autonomous vehicles, paving the way for much of the current work in autonomous cars and drones [38]. Other notable DARPA programs since 2007 include:

- Cognitive Technology Threat Warning System [39]
- Human Assisted Neural Devices [40]
- Autonomous Real-time Ground Ubiquitous Surveillance-Imaging System [41]
- Urban Reasoning and Geospatial Exploitation Technology [42]
- Adaptive Radar Countermeasures [43]
- Behavioral Learning or Adaptive Electronic Warfare [44]
- Mind's Eye [45]
- Probabilistic Programming for Advancing Machine Learning [46]
- Explainable Artificial Intelligence [47]

AI is also being integrated into existing DARPA programs. In 2016, DARPA launched the *Sea Hunter*, which had been developed through its Anti-submarine Warfare Continuous Trail Unmanned Vessel program. The *Sea Hunter* is the first unmanned ship that can perform most naval warship functions and uses AI to operate with a minimal crew. Sea trials are underway, and it is expected to join the fleet in 2018 [48].

Beyond DARPA, other military agencies have created robust AI-centered programs. In 2010, the Air Force Research Laboratory (AFRL) created Synchronized Net-Enabled Multi-INT Exploitation, a project aimed at developing “advanced intelligence-collection and rapid-response capability” to assist U.S. forces with attacking the enemy at its weakest points [49]. The project emphasized machine-to-machine intelligence communications and cooperation and utilized AI and other techniques to more effectively handle the massive amounts of data collected both on the battlefield and through intelligence. In 2016, AFRL ran simulations between a drone and an experienced fighter pilot, and the pilot was defeated in every scenario. The drone employed software called ALPHA, which uses fuzzy logic in its AI algorithms to mimic a human’s decision cycle [50]. The Space and Naval Warfare Systems Center Intelligent Technologies for Decision Making program was initiated to explore uses of AI for ship’s navigation, tactical analysis, sensors, and weapons [51].

The use of AI in intelligence gathering is another area of robust research. While related to and often overlapping with military operations, most of this work has been coordinated through the Intelligence Advanced Research Projects Activity (IARPA). Similar to DARPA, IARPA invests in high-risk/high-reward research but in regards to intelligence activities. Notable IARPA AI research programs from 2010 to the present include:

- Integrated Cognitive-Neuroscience Architectures for Understanding Sensemaking [52]
- Knowledge Representation in Neural Systems [53]
- Strengthening Human Adaptive Reasoning and Problem-Solving [54]
- Machine Intelligence from Cortical Networks [55]

Homeland security is another area in which AI research is taking place. In 2008, a joint project between DHS and AFRL was launched. The Distributed Environment for Critical Infrastructure Decision-making Exercises uses AI to simulate players in disaster scenarios [56]. In 2012, U.S. Customs and Border Protection piloted a two-month study at the border crossing in Nogales, Arizona, using the Automated Virtual Agent for Truth Assessments in Real-Time (AVATAR), which is a virtual agent that scans an individual's identification then uses a combination of a questionnaire and sensors to create a risk assessment. The platform uses hundreds of different psychophysiological and behavioral cues to determine a person's credibility. Anyone who is flagged in this initial screening is directed to undergo a secondary screening with a human agent [57]. AVATAR is also being tested by the Canadian Border Services Agency [58], and is expected to play a key part in developing border crossings of the future.

AI played a central role in the DoD's Third Offset Strategy, which was unveiled in 2014 [59]. An offset strategy is long-term, competitive strategy aimed on generating and sustaining a strategic advantage between defense establishments [60]. The First Offset was developed in the 1950s as a response to Cold War tensions with the Soviet Union and focused on nuclear armament and deterrence. The Second Offset was developed in the 1970s and 1980s and focused on precision-guided weapons systems. The Third Offset focused on developing technological advancements in machine learning and autonomy and included five components:

1. Learning Systems
2. Human-machine Collaboration
3. Human-machine Combat Teaming
4. Assisted Human Operations
5. Network-enabled, Cyber-hardened Weapons [61]

The Third Offset Strategy was rolled out to counter similar technological advances in weapons systems and techniques from other nations. In December 2015, during a speech hosted by the Center for a New American Security, Deputy Secretary of Defense Bob Work commented on the Third Offset, stating, "Now, we believe, strongly, that humans should be the only ones to decide when to use lethal force. But when you're under attack, especially at machine speeds, we want to have a machine that can protect us [62]." Secretary Work remarked in October 2016 at the Center for Strategic and International Studies that "the pacing competitors—not adversaries—are Russia and China, because they're developing advanced capabilities that potentially worry us [63]." Secretary Work has also stated that "over time we'll go to exercises where all five of the pieces are operating together [61]."

As the prospect of increased autonomy and use of AI in weapons systems has become more likely, the DoD realized the importance of establishing procedures for dealing with these types of weapons. In 2012, DoD Directive 3000.09 (Autonomy in Weapons Systems) was issued. In general, the directive guides the development of autonomous and semi-autonomous systems. In particular, it expresses the intent to ensure that a human being always has the final authorization regarding the execution of commands that would fire at a target [64]. The directive also prompted a review of autonomous systems within the DoD by

the Defense Science Board in 2012 [65], and later the Summer Study on Autonomy, which concluded that “while difficult to quantify, ...autonomy—fueled by advances in artificial intelligence—has attained a ‘tipping point’ in value [66].” The announcement of the Third Offset Strategy, recommendations from the DoD and other organizations, and the rapid rise of AI research in industry prompted action for a unified national AI strategy. In 2016, the Obama administration established the Subcommittee on Machine Learning and AI and released the *National Artificial Intelligence Research and Development Strategic Plan* [1]. It is rather difficult to pinpoint all of the areas within the DoD in which AI research takes place, but DARPA, the U.S. Army Research Laboratory, ONR, and AFRL are certainly leading the effort, and Deputy Secretary Work has stated that the DoD is expected to spend \$12–15 billion on AI research in the fiscal year 2017 [19].

Methodology and Structure of the Report

It is assumed that the reader will have a basic understanding of AI. Since most DoD AI research is applied research, this report focuses much more on applications of AI and not fundamentals of programming, mathematics, or computer science needed for AI. The intent is to illustrate what is currently possible across defense, homeland security, and intelligence applications and not to focus on nuanced developments in the field itself. Developments over the past three years and what is possible within the next 18 to 24 months have been the most scrutinized aspects of AI research for this report. HDIAC reached out to subject matter experts for their opinions and knowledge of AI as it applies to HDIAC’s eight focus areas. Because this report focuses on applications, it is divided into sections based on HDIAC focus areas.

References

1. Executive Office of the President. National Science and Technology Council Committee on Technology. (2016). Preparing for the future of artificial intelligence. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf
2. Corrin, A. (2017, December 6). DoD spent \$7.4 billion on big data, AI and the cloud last year. Is that enough? Retrieved from <https://www.c4isrnet.com/it-networks/2017/12/06/dods-leaning-in-on-artificial-intelligence-will-it-be-enough/>
3. China State Council. (2017). State Council Notice on the Issuance of the NextGeneration Artificial Intelligence Development Plan (translated). Retrieved from <https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems/>
4. CIFAR. (n.d.). Pan-Canadian artificial intelligence strategy overview. Retrieved from <https://www.cifar.ca/assets/pan-canadian-artificial-intelligence-strategy-overview/>
5. American Psychological Association. (n.d.). Intelligence. Retrieved from <http://www.apa.org/topics/intelligence/>
6. Merriam-Webster. (n.d.). Intelligence [Definition 1A]. Retrieved from <https://www.merriam-webster.com/dictionary/intelligence>
7. Gottfredson, L. S. (1997). Mainstream science on intelligence: An editorial with 52 signatories, history, and bibliography. *Intelligence*, 24(1), 13-23. doi:10.1016/S0160-2896(97)90011-8
8. Tirri, K., & Nokelainen, P. (2011). *Measuring multiple intelligences and moral sensitivities in education* (5th ed.). Rotterdam, The Netherlands: Sense Publishers.
9. Poole, D. L., Mackworth, A. K., & Goebel, R. G. (1998). *Computational intelligence: A logical approach*. New York, NY: Oxford University Press.
10. Russell, S., & Norvig, P. (2014). *Artificial intelligence: A modern approach* (3rd ed.). Harlow, United Kingdom: Pearson Education Limited.
11. Stanford Encyclopedia of Philosophy. (2013, August 28). Classical logic. Retrieved from <https://plato.stanford.edu/entries/logic-classical/Logic>
12. Weed, L. (2011, November/December). Philosophy of mind: An overview. *Philosophy Now*, (87). Retrieved from https://philosophynow.org/issues/87/Philosophy_of_Mind_An_Overview
13. Merriam-Webster. (n.d.). Psychology. Retrieved from <https://www.merriam-webster.com/dictionary/psychology>
14. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., . . . Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526. doi:10.1073/pnas.1611835114
15. Lighthill, J. (1973). *Artificial intelligence: A general survey*. Retrieved from http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/contents.htm
16. National Research Council. (1999). *Funding a revolution: Government support for computing research*. Washington, DC: The National Academies Press. doi:10.17226/6323.
17. Tilghman, P. (2017, July 13). Conversation with DARPA Program Manager [Telephone interview].

18. Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., & Trench, M. (2017). Artificial intelligence: The next digital frontier. McKinsey Global Institute. Retrieved from <https://www.mckinsey.com/~media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx>
19. Shalal, A. (2015, December 14). The Pentagon wants at least \$12 billion to fund AI weapon technology in 2017. Business Insider. Retrieved from <http://www.businessinsider.com/the-pentagon-wants-at-least-12-billion-to-fund-ai-weapon-technology-in-2017-2015-12>
20. Rejewski, M. (1981). How Polish mathematicians deciphered the Enigma. *IEEE Anal. of the History of Computing*, 3(3), 213-234. Retrieved January 28, 2018 from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.692.9386&rep=rep1&type=pdf>
21. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460. Retrieved from <http://www.jstor.org/stable/2251299>
22. McCarthy, J., Minsky, M.L., Rochester, N., & Shannon, C.E. (2006). A proposal for the Dartmouth Summer Research Project on Artificial Intelligence. *AI Magazine*, 27(4), 12-14. Retrieved from <https://www.aaai.org/ojs/index.php/aimagazine/article/view/1904>
23. Copeland, B. J. (1993). *Artificial Intelligence: A philosophical introduction*. Oxford: Wiley-Blackwell.
24. *Mark I Perceptron at the Cornell Aeronautical Laboratory* [Photograph found in Cornell University News Services Records 4-3-15, Division of Rare and Manuscript Collections, Cornell University Library]. (n.d.). Retrieved from <https://digital.library.cornell.edu/catalog/ss:550351>
25. Bishop, C. M. (2006). *Pattern recognition and machine learning*. New York, NY: Springer-Verlag.
26. Howe, J. (2007). The nature of artificial intelligence. Retrieved from <http://www.inf.ed.ac.uk/about/Alhistory.html>
27. National Research Council. (1966). *Language and machines: Computers in translation and linguistics*. Washington, DC: The National Academies Press. doi:10.17226/9547
28. Ireland, C. (2012, September 13). Alan Turing at 100. *The Harvard Gazette*. Retrieved from <http://news.harvard.edu/gazette/story/2012/09/alan-turing-at-100/>
29. Colby, K.M., Watt, J.B., & Gilbert J.P. (1966). A computer method of psychotherapy: preliminary communication. *The Journal of Nervous and Mental Disease*. 142(2), 148-152. Retrieved from https://journals.lww.com/jonmd/Citation/1966/02000/A_Computer_Method_of_Psychotherapy__Preliminary.5.aspx
30. Lighthill, J. (1973). *Artificial intelligence: A general survey*. Retrieved from http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/contents.htm
31. Laitinen, H.A. (1970). Reverberations from the Mansfield Amendment. *Anal. Chem.* 42 (7), 689-689. doi: 10.1021/ac60289a600
32. Williams, P. (1976). Whatever happened to the Mansfield amendment? *Survival: Global Politics and Strategy*, 18(4), 146-153. doi:10.1080/00396337608441623

33. Korolkov, A. (2014, July 2). New combat robot is Russian army's very own deadly WALL-E. Retrieved from https://www.rbth.com/defence/2014/07/02/new_combat_robot_is_russian_armys_very_own_deadly_wall-e_37871.html
34. Encyclopedia Britannica. (2008). MYCIN. Retrieved from <https://www.britannica.com/technology/MYCIN>
35. Grosz, B. & Davis, R. (n.d.) A report to ARPA on twenty-first century intelligent systems. American Association for Artificial Intelligence. Retrieved from <https://aaai.org/Library/Reports/arpa-report.php>
36. Krauthammer, C. (1997, May 26). Be afraid: The meaning of Deep Blue's victory. The Weekly Standard. Retrieved from <http://www.weeklystandard.com/be-afraid/article/9802>
37. SRI International. (2003). Artificial intelligence: CALO. Retrieved from <https://www.sri.com/work/timeline-innovation/timeline.php?timeline=computing-digital#!&innovation=artificial-intelligence-calo>
38. DARPA. (2013, March 13). The DARPA grand challenge: Ten years later. Retrieved from <https://www.darpa.mil/news-events/2014-03-13>
39. DARPA. (2007). Broad agency announcement (BAA) 07-25, Cognitive technology threat warning systems (CT2WS). Retrieved from <https://web.archive.org/web/20080204203721/http://www.darpa.mil/baa/BA07-25.html>
40. Miranda, R. A., Casebeer, W. D., Hein, A. M., Judy, J. W., Krotkov, E. P., Laabs, T. L., . . . Ling, G. S. (2015). DARPA-funded efforts in the development of novel brain-computer interface technologies. *Journal of Neuroscience Methods*, 244, 52-67. doi:10.1016/j.jneumeth.2014.07.019
41. DARPA. (n.d.). Autonomous Real-time Ground Ubiquitous Surveillance - Infrared (ARGUS-IR). Retrieved from <https://www.darpa.mil/program/autonomous-real-time-ground-ubiquitous-surveillance-infrared>
42. Rutherford, M. (2009, August 21). DARPA 3D reasoning engine to identify urban threats. Retrieved from <https://www.cnet.com/news/darpa-3d-reasoning-engine-to-identify-urban-threats/>
43. DARPA. (n.d.). Adaptive Radar Countermeasures (ARC). Retrieved from <https://www.darpa.mil/program/adaptive-radar-countermeasures>
44. DARPA. (n.d.). Behavioral Learning for Adaptive Electronic Warfare (BLADE) Retrieved from <https://www.darpa.mil/program/behavioral-learning-for-adaptive-electronic-warfare>
45. Keller, J. (2010, May 1). Artificial intelligence and machine vision for unmanned vehicle smart cameras is aim of DARPA Mind's Eye program. Retrieved from <http://www.militaryaerospace.com/articles/print/volume-21/issue-5/electro-optics-watch0/artificial-intelligence.html>
46. DARPA. (n.d.). Probabilistic Programming for Advancing Machine Learning (PPAML). Retrieved from <https://www.darpa.mil/program/probabilistic-programming-for-advancing-machine-learning>
47. DARPA. (n.d.). Explainable Artificial Intelligence (XAI). Retrieved from <https://www.darpa.mil/program/explainable-artificial-intelligence>
48. Ghose, T. (2016, April 8). 'Sea Hunter': World's first unmanned ship stalks subs. Retrieved from <https://www.livescience.com/54354-darpa-launches-unmanned-submarine-hunter.html>

49. Keller, J. (2010, March 1). Air Force to use artificial intelligence and other advanced data processing to hit enemy where it hurts. *Military and Air Force Electronics*, 21(3). Retrieved from <http://www.militaryaerospace.com/articles/print/volume-21/issue-3/news/news/air-force-to-use-artificial-intelligence-and-other-advanced-data-processing-to-hit-the-enemy-where-it-hurts.html>
50. Clark, C. (2016, August 8). Artificial intelligence drone defeats fighter pilot: The future? Breaking Defense. Retrieved from <http://breakingdefense.com/2016/08/artificial-intelligence-drone-defeats-fighter-pilot-the-future/>
51. Keller, J. (2012, June 22). Navy eyes artificial intelligence to automate decision making on ships, sensors, and weapons. *Military & Aerospace Electronics*. Retrieved from <http://www.militaryaerospace.com/articles/2012/06/artificial-intelligence-navy.html>
52. Kevin Burns, K., Bonaceto, C., Fine, M., & Oertel, C. (2014). Integrated cognitive-neuroscience architectures for understanding sensemaking (ICArUS): Overview of test and evaluation materials. MITRE. Retrieved from <http://www.dtic.mil/dtic/tr/fulltext/u2/a624922.pdf>
53. IARPA. (n.d.). Knowledge Representation in Neural Systems (KRNS). Retrieved from <https://www.iarpa.gov/index.php/research-programs/krns>
54. IARPA. (n.d.). Strengthening Human Adaptive Reasoning and Problem-Solving (SHARP). Retrieved from <https://www.iarpa.gov/index.php/research-programs/sharp>
55. IARPA. (n.d.). Machine Intelligence from Cortical Networks (MICrONS). Retrieved from <https://www.iarpa.gov/index.php/research-programs/microns/microns-baa>
56. U.S. Department of Homeland Security. (n.d.). Distributed environment for critical infrastructure decision-making exercises. Retrieved from <https://www.dhs.gov/science-and-technology/csd-decide>
57. U.S. Department of Homeland Security. (n.d.). Rapid screening tool: The AVATAR . Retrieved from <https://www.dhs.gov/sites/default/files/publications/Rapid%20Screening%20Tool-NCBSI-AVATAR-Jan2014.pdf>
58. Finch, S. (2016, December 28). The lie-detecting security kiosk of the future. Phys.org. Retrieved from <https://phys.org/news/2016-12-lie-detecting-kiosk-future.html>
59. Martinage, R. (2014). Toward a new offset strategy exploiting U.S. long-term advantages to restore U.S. global power projection capability. Center for Strategic and Budgetary Assessments. Retrieved from <http://csbaonline.org/uploads/documents/Offset-Strategy-Web.pdf>
60. Lange, K. (2016, March 30). 3rd Offset Strategy 101: What it is, what the tech focuses are. DoDLive. Retrieved from <http://www.dodlive.mil/2016/03/30/3rd-offset-strategy-101-what-it-is-what-the-tech-focuses-are/>
61. Pellerin, C. (2015, November 13). Learning systems, autonomy and human-machine teaming. Armed with Science. Retrieved from <http://science.dodlive.mil/2015/11/13/learning-systems-autonomy-and-human-machine-teaming/>
62. Work, B. (2015, December 14). *Deputy Secretary of Defense Speech*. Speech presented at CNAS Defense Forum in Washington, DC.

- Retrieved from <https://www.defense.gov/News/Speeches/Speech-View/Article/634214/cnas-defense-forum/>
63. Pellerin, C. (2016, October 31). Deputy Secretary: Third Offset Strategy bolsters America's military deterrence. U.S. Department of Defense. Retrieved from <http://www.defense.gov/News/Article/Article/991434/deputy-secretary-third-offset-strategy-bolsters-americas-military-deterrence?source=GovDelivery>
 64. DoD Directive 3000.09. (Updated 2017, May 9). Autonomy in Weapons Systems. Retrieved from <http://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf>
 65. Defense Science Board. (2012, July). *DSB task force report: The role of autonomy in DoD systems*. Washington, DC. Retrieved from <https://www.hsdl.org/?view&did=722318>
 66. Defense Science Board. (2016, June). *Defense Science Board summer study on autonomy*. Washington, DC. Retrieved from <https://www.hsdl.org/?abstract&did=794641>

2

Use of Artificial Intelligence in Enabling Renewable Energy

*Sue Ellen Haupt, National Center for Atmospheric Research
David John Gagne, National Center for Atmospheric Research*

Introduction

Growth of Renewable Energy

The growth in deployment of renewable energy has been accelerating over the past decade. By 2016, new installations of renewables worldwide were 161 gigawatts (GW) [1] for a total global capacity of 2,006 GW [2]. U.S. installed capacity of wind energy at the end of 2016 was 82 GW, which included 8.2 GW of new wind capacity, representing 41 percent of all new capacity in 2015 [3]. U.S. solar capacity reached 44.7 GW, with an average annual growth rate of 68 percent over the past decade [4], and, in 2016, utility-scale solar comprised 72 percent of those capacity installations [4]. Investment in new renewable capacity worldwide in 2016 was approximately double that in fossil fuels capacity for the fifth consecutive year [1]. However, much of the worldwide growth in renewables has been in Asia, with China becoming the world leader in the expansion of renewable energy. It is estimated that by 2021, China will have more than one-third of the global capacity of solar photovoltaic (PV) and onshore wind capacity [5]. India is also showing more interest in installation of renewables and its solar PV capacity is expected to grow by a factor of eight [5]. Aid organizations are working with developing countries to stimulate their deployment of renewables as well.

This growth in renewables has been spurred by cost reductions, the drive to reduce emissions of CO₂ and other pollutants, and the aim of countries to achieve energy independence. In 2016, CO₂ emissions from fossil fuels stabilized globally for the third consecutive year despite global economic growth of 3 percent and increased energy demand [1]. Policies in many countries have evolved to become more favorable to renewables, which has helped enhance deployment [5].

Growth of Forecasting for Renewables and Necessity of Using AI

Because renewable energy uses natural resources (e.g., wind, solar, hydro) as its fuel, that fuel supply is inherently variable. Utilities and independent systems operators (ISOs) deal with that variability by incorporating forecasts [6]. These forecasts are as critical to planning their operations as forecasts of hourly energy load and the knowledge of coal, oil, gas, and nuclear unit availability. Lew et al. [7] discuss the challenge of solar power integration, showing that the variability of power output was greater with high penetrations of solar than with high penetrations of wind. The response speed (ramp rate and start time), response duration, frequency of use (continuously or only during rare events), direction of use (up or down), and type of control characterize a utility company's operating reserves [8]. And these operating reserves can be managed with accurate solar forecasts. Curtright and Apt [9] have shown that the cost of energy can be strategically minimized with knowledge of the short- and long-term variations. Some countries manage very high peaks in energy production from renewables,

including energy peaks of 140 percent for Denmark and 86.3 percent for Germany in 2016 [1] by maintaining high-quality forecasts and efficient operating procedures. In the U.S., Xcel Energy (the American utility with the highest percentage of wind capacity) calculates that accurate forecasts of its wind power saves its customers roughly \$10 million per year and also substantially reduces emissions of CO₂ and other pollutants [10]. These forecast systems leverage both physical models and those based on AI. AI forecast systems based on machine learning models that utilize a blend of weather observations and numerical weather prediction (NWP) model output have consistently produced the most accurate power forecasts and have been widely adopted by electric utilities and power forecasting companies [11-13].

Similarly, military energy needs must be balanced, both for stationary microgrids and for supplying energy to mobile troops. Just as the users must be aware of the fluctuations in load, they must also be aware of fluctuations of the renewable resources on various scales. This section of the SOAR will focus on the traditional utility and ISO use and then draw parallels for applications to the DoD.

Military Applications of Renewables

As the nation's largest energy user, DoD has a goal of managing its energy resources judiciously. A 2014 DoD Directive [14] clarifies policies to manage energy so as to enhance military capacity, improve energy security, and mitigate costs. To this end, the DoD seeks to improve energy performance of its systems, diversify energy supplies to include renewable energy, and include energy analyses in its planning [14]. The DoD signed a Memorandum of Understanding (MOU) with the U.S. Department of Energy to establish energy boards and councils [15]. The 2010 MOU states:

Energy efficiency can serve as a force multiplier, increasing the range and endurance of forces in the field while reducing the number of combat forces diverted to protect energy supply lines, as well as reducing long-term energy costs. DoD is also increasing its use of renewable energy supplies and reducing energy demand to improve energy security and operational effectiveness, reduce greenhouse gas (GHG) emissions in support of U.S. climate change initiatives, and protect the DoD from energy price fluctuations. Solving military challenges through innovation has the potential to yield spin-off technologies that benefit the civilian community as well [15].

In its *2016 Operational Energy Strategy* [16], the DoD listed a specific goal to develop energy supplies to reduce risk. As part of meeting that goal, DoD is pursuing renewable energy opportunities, including harvesting energy from the environment, particularly solar energy. To that end, the Army Research Laboratory has assessed the integration of meteorological forecasting into the use of solar power, including exploring solar power in microgrids [17,18].

By incorporating more renewable energy resources into its infrastructure, DoD can increase its resilience to accidental and adversarial disruptions of traditional electricity sources and reduce both the financial and personnel costs of maintaining a fossil fuel supply chain [19]. Military bases in both the U.S. and those overseas are dependent on a mixture of electricity supplied by outside sources

and diesel generators in order to maintain continuous operations, and the transmission lines and power plants outside military bases are vulnerable to disruption. Additionally, diesel generators require constant resupply of fuel by convoys, and in Iraq and Afghanistan, these convoys were frequently attacked, resulting in hundreds of casualties [20]. On-base renewable energy sources are better protected from attack and can offer more resilience, even when other parts of the grid are disrupted.

Wind and solar energy resources produce power amounts that fluctuate based on weather conditions, so successful integration of these sources into a power grid requires accurate forecasts of electricity supply and demand. Although improving battery storage technology has been a major focus in making renewable energy viable for military applications [21], having storage available does not remove the need for forecasts. While storage can help reduce the impacts of sudden supply fluctuations, grid operators still need renewable energy power forecasts to allocate power resources efficiently.

The remainder of this section discusses how AI methods are currently implemented in providing those forecasts at different forecast time horizons for civilian electric grids and how these methods can be adapted for military operations. After discussion of the current state-of-the-art systems, promising future directions are highlighted.

Forecasting Renewable Energy and AI

Overview

Forecasting renewable energy has been widely documented in various books [22,23] and journal papers. Many of the techniques employed include AI as described here and must be designed for the time scale of the fluctuation in the resource. Fluctuations of the renewable resource are separated into very short range (hours ahead), medium range (day ahead to week ahead), and longer range (weeks to seasons and decades).

At the short range of hours ahead, the issue is balancing the load in the short range, typically performed by the ISO grid operators. Because the load fluctuates, as do the renewable resources, operators rely on having real-time short-range forecasts for situational awareness (often called nowcasts) of the renewable resources, in this case of wind and solar power. In grid operating rooms, operators have displays of the conventional units currently operating as well as displays of real-time forecasts of the expected renewable resources. Grid operators function on time scales of five- to 15-minute forecasts and expect information at that resolution for the next 3 to 6 hours.

The medium-range forecasts focus on the day ahead for unit allocations and energy trading. Every morning, utilities bid expected power production for each unit according to each unit's marginal price, which largely depends on fuel price and operating parameters. For wind and solar power, the fuel is free, meaning that the marginal prices are quite inexpensive and those units are typically allocated at their maximum expected amounts. Thus, an accurate day-ahead forecast is required to know the amount of wind and solar available. In addition, trading between utilities is based on optimizing the marginal prices and utilizing those

same forecasts. Hourly forecasts are required for the following day (often around 6 a.m. of the current day, implying a forecast for a lead time of 18 to 42 hours) so unit allocations can be applied for hourly timeframes. Because utilities operate primarily during normal business hours, a three-day forecast is needed for weekend periods. And when holidays lengthen the weekend, the medium-range forecast becomes four or five days.

Longer-term forecasts are needed for longer-term fuel optimization and maintaining equipment, including transmission and distribution lines. In addition, weather information helps energy companies prepare for extreme events (e.g., heat waves, cold snaps, icing, and very extreme weather such as hurricanes). The longest-range forecasts (decades) are needed for resource assessment to plan the best locations for renewable power plants. Not only do renewable developers need to know the mean expected output of a potential plant, but they also need to know likely fluctuations on scales ranging from the very short through interannual and climatic scales. Thus, forecasting for these uses requires long-term observations and AI capabilities to better predict expected renewable output.

Use of AI on the Very Short Range (Hours)

The shortest ranges (0–3 or 6 hours) is where AI outshines most other techniques for forecasting renewable energy production. As such, AI, along with other statistical learning methods, has been the mainstay of this prediction range. The most popular method used is the artificial neural network [24-37]. Pedro and Coimbra [25] showed that for one- to two-hour solar power predictions, an artificial neural network model for a time series outperformed persistence, AutoRegressive Integrated Moving Average (ARIMA), and k-Nearest Neighbors models.

Reikard, et al. [38] found that ARIMA models with dynamic coefficients are good at this for short range but blending NWP models with dynamic coefficients is even better. Lopez, et al. [39] accomplished a sensitivity study of the most relevant predictors for estimating Direct Normal Irradiance with a Bayesian Artificial Neural Network (ANN) method and chose to use the clearness index and the relative air mass. The large number of tunable parameters requires a large quantity of training data to prevent overfitting or loss of skill.

Some studies have found value in applying clustering techniques in advance of training AI algorithms as discussed below. The philosophy behind this approach is that weather occurs in regimes, and separating the data into specific regimes may provide clusters of training data that behave similarly. Thus, one might expect that the regime-dependent AI models produce better predictions than possible without the clustering [40]. Hinkelman [41] showed that solar irradiance variability differs among satellite data-derived cloud types, which suggests that this approach is appropriate for irradiance forecasting.

A review of various methodologies for both supervised and unsupervised cloud classification is provided by Tapakis and Charalambides [42]. The supervised techniques classify regimes based on available training datasets and arithmetic complexity of the technique, while the unsupervised techniques classify based on the pixels of an image. McCandless, et al. [43] and Morf [44] demonstrated a one-step stochastic prediction process for cloud cover or clearness index. K-means clustering was used by Zagouras et al. [45] to identify regimes based on step-

changes of the average daily clear sky index in the San Diego, California region. Mellit et al. [46] trained separate models for three simple classifications of daily total solar irradiance into clear, partly cloudy, and cloudy regimes and found improvement over a single all-sky model, particularly for the cloudy days. McCandless et al. used a k-means clustering algorithm to identify regimes before applying an ANN [47]. That work used data regarding surface weather and irradiance observations in addition to direct observation of irradiance values.

At times, it is necessary to have a very accurate, very short-range forecast for the first 15 to 30 minutes. In this case, instruments at the site provide situational awareness and the ability to forecast based on observations in the immediate vicinity. One instrument that has been quite useful for these scales is the sky imager, which consists of a camera that either looks upward or, in some cases, is pointed at a mirror that looks upward. These instruments provide a real-time view of the sky with its passing clouds. Using successive frames as input to cloud-based advection techniques has proven to be a successful way to optimize prediction in the first 15 to 30 minutes [37,48-51]. Some groups have deployed multiple imagers in the same vicinity to expand the range of view, as well as to allow discernment of the cloud base and height data, allowing motion vectors to vary with height [50,52]. These methods tune the motion vector algorithms with AI methods, such as support vector machines [52].

Satellites afford remote sensing data that also provide situational awareness and the capability to process successive frames to predict cloud motion [53-55]. Some researchers have combined AI methods with satellite-based cloud-vector motion methods. Marquez et al. [32] showed a five percent and 25 percent reduction in Root-Mean-Squared Error (RMSE) compared to that of persistence by integrating processed satellite images as input into ANNs to predict global horizontal irradiance (GHI) from 30 minutes to 120 minutes. McCandless et al. [56] integrated satellite-based cloud classification information into regime-dependent neural networks.

Figure 1 demonstrates the improvement of the regime-dependent techniques over a “smart” persistence forecast (persistence in cloud cover, allowing for changes in the solar angle) beyond 15 minutes. Beyond the 15-minute lead-time, most of the regime-dependent ANN-trained methods show between 10 percent and 28 percent improvement over the clearness index persistence method. Using the cloud classification directly from the satellite classification does not improve performance on average. The regimes using the satellite information but classified via k-means clustering showed 21.0 percent, 26.4 percent, and 27.4 percent improvement over the clearness index persistence at the one-hour, two-hour, and three-hour forecast lead times, respectively. This method is the best overall solution. In general, this work indicates that cloud regime classification that includes satellite information makes a considerable positive impact on the overall performance of the models [57].

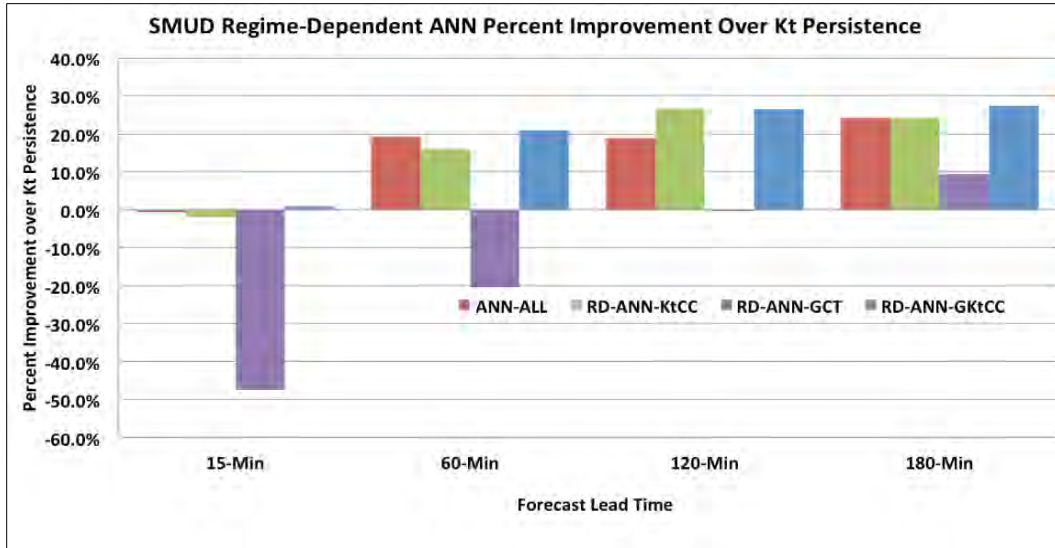


Figure 1. Percent improvement of Regime-Dependent ANNs over Smart Persistence near Sacramento, CA. ANN-All (red) trains an ANN on all data without separation into regimes, RD-ANN-GCT (purple) uses regimes derived from the GOES satellite data, RD-ANN-KtCC used cloud cover regimes derived via k-means clustering, and RD-ANN-GKtCC (blue) uses cloud cover regimes that include satellite-based information [57]. ©American Meteorological Society. (Used with permission)

Finally, having nearby observations of related weather conditions can be leveraged to produce forecasts at the site of the renewable energy resource. When such observations are available upstream of renewable energy sites, it is possible to build “expert systems” that make use of this information together with average speed and direction of system movement. Kosovic et al. [10] describe application of one such system to predict wind energy.

Use of AI for Medium Range Forecasting

The primary forecasting tool for the medium range is NWP, which is based on the Navier-Stokes equations of fluid motion combined with the forcing physics that describe radiative forcing, land-surface processes, cloud-physics processes, boundary-layer turbulence, and more. Forecasters have found they can improve the accuracy of the NWP forecasts by calibrating and smartly blending via AI and statistical learning methods [58-62]. Recently, some NWP models have been tailored specifically to forecast renewable energy and its impacts, including the Weather Research and Forecasting (WRF) model [63-65].

The most basic calibrating method is Model Output Statistics (MOS), first demonstrated by Glahn and Lawry [66], which uses multilinear statistical regression to calibrate the forecasts. Although this method requires a sufficiently long time series to train (one to two years of data), it is in use in forecasting centers throughout the world for a variety of forecast variables. Most of the forecast equations are regionalized to increase the sample size [67].

AI methods have been utilized for day- to week-ahead lead times to generate bias-corrected forecasts of renewable energy resources, such as wind speed and solar irradiance, and estimates of the power produced from those resources. AI models applied within this time frame generally use a mixture of atmospheric variables from one or more NWP models as input, along with indicators of time of day,

season, and terrain information. Each NWP model exhibits different systematic biases, so the optimal deterministic forecast is found by first bias correcting the forecast from each individual NWP model and then blending the bias corrected forecasts. Every major type of machine learning model has been evaluated for use in wind and solar power forecasting at short- and moderate-range time frames.

An example of a blending system used by several forecasting companies is the National Center for Atmospheric Research’s (NCAR) Dynamic Integrated foreCast system [68], DICAST®, which produces these optimal blended forecasts for both wind [58,69] and solar [61,70] resources. DICAST blends output from multiple NWP models using a two-step process. First, each model is calibrated using a dynamic version of MOS. Second, the models are blended via optimizing the weights using a dynamic weighting algorithm. DICAST is typically applied separately at each site of interest and for each forecast lead time. Weights are updated frequently. An advantage of DICAST is that it can optimize forecasts with less data than many other AI algorithms, providing forecasts using as little as 30 days of historical data and optimized at around 90 days of data. DICAST has been shown to decrease the error beyond the best model by 10 to 15 percent [58,71]. Figure 2 shows the performance of DICAST as the SunCast (represented by the yellow line) forecast for day-ahead forecasts. The mean absolute error of the SunCast system is consistently better than that of any of the component models.

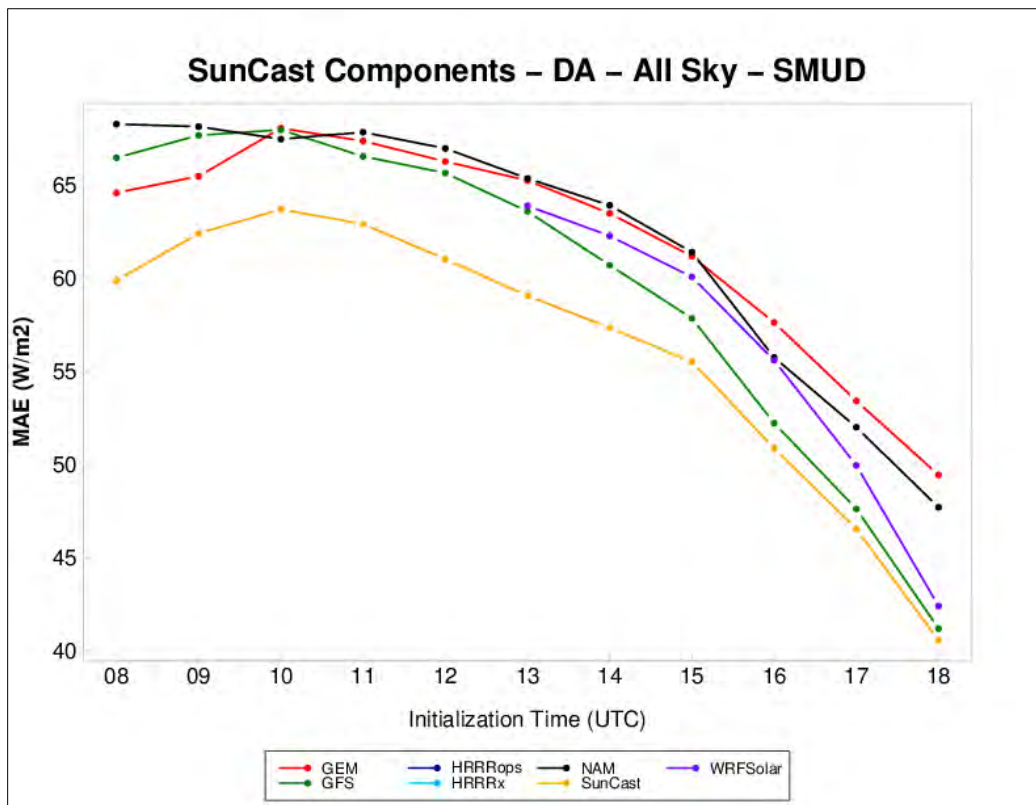


Figure 2. Mean Absolute Error in W m-2 for all Day-Ahead forecasts from DICAST® components and the DICAST blended Sun4Cast® system near Sacramento, California for all sky conditions. The blended system has substantially lower error than any component model [70]. ©American Meteorological Society. (Used with permission)

While primary operating bases are more likely to have the observational infrastructure to allow for the training of AI systems dedicated to those bases, forward operating bases and warfighters operating portable solar systems will not be able to utilize a solar forecasting model trained to their location. Instead, they could use gridded renewable energy forecasting models that are trained on all available data in the area and are applied at the location of interest. Gagne et al. [72] evaluated different gridded solar irradiance forecasting methods in Oklahoma and found that gradient-boosted regression trained over multiple sites and applied directly at the site of interest resulted in the lowest mean absolute error. Forecast errors increased for all methods when conditions were partly cloudy and at sites with higher amounts of cloudiness. The largest errors occurred in the spring, and the smallest errors occurred in the winter. Figure 3 demonstrates performance of that model.

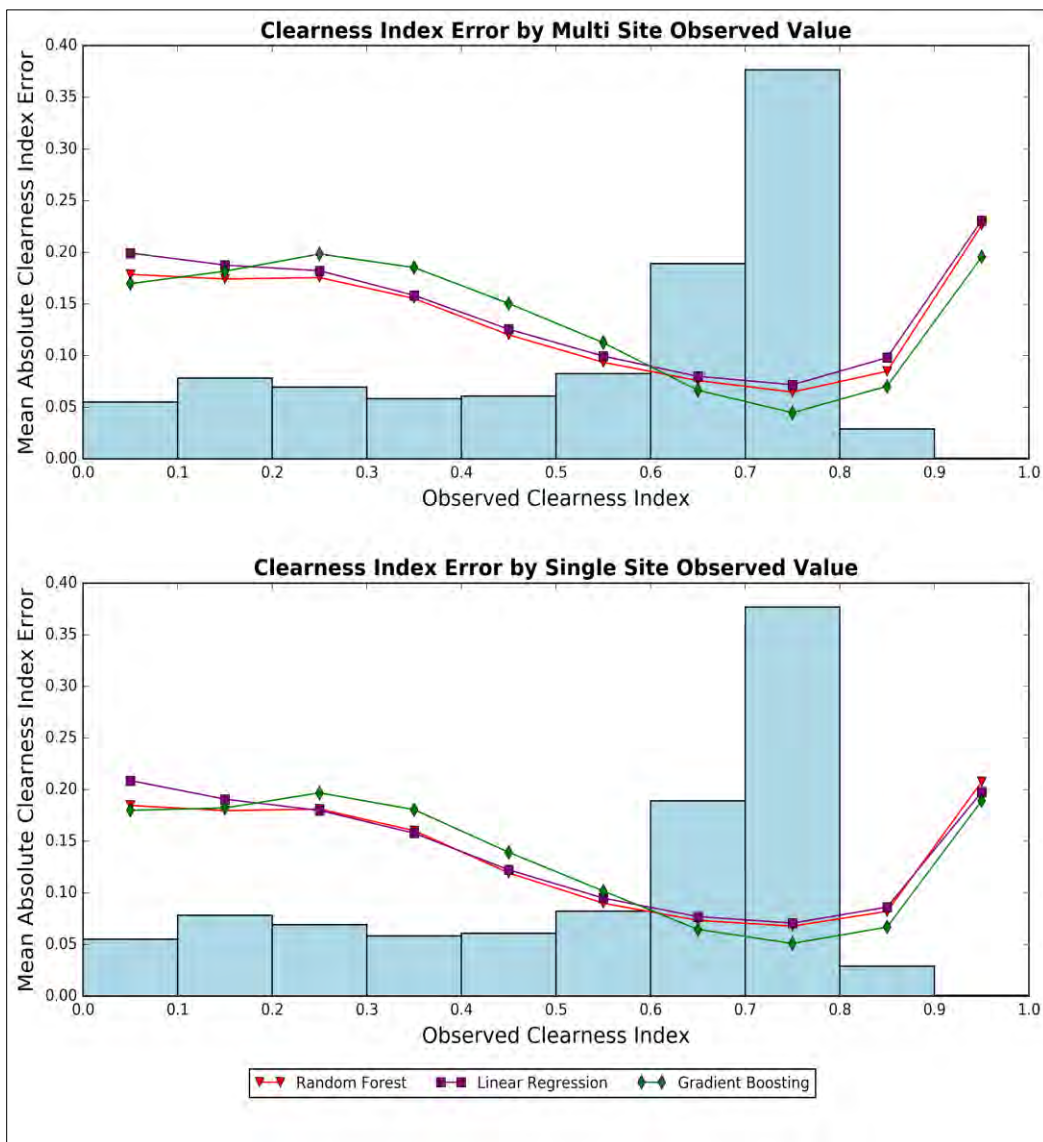


Figure 3. Mean absolute error of clearness index forecasts in Oklahoma grouped by the observed clearness index value for three types of machine learning models. The histogram in blue displays the observed relative frequency of clearness index values [72].

McGovern et al. [73] conducted an international forecasting contest to predict total daily solar radiation at sites across Oklahoma using an ensemble of global numerical weather prediction model output. Many machine learning methods were tested on this dataset, and the top performers all used gradient-boosting regression in their systems. The key differences among the methods resulted from how the model forecasts were interpolated from the model grid points to the forecast sites. Combining multiple interpolation approaches improved the predictions because the different approaches brought different kinds of information to the sites.

Some Issues for Distributed Solar Forecasting

One confounding issue for forecasting solar power is that much of the deployment is behind the meter, making it difficult to separate production from this Distributed Photovoltaic (DPV) resource from reductions in load at that site. There are several approaches to addressing this problem, primarily a bottom-up or a top-down approach [12]. The bottom-up approach requires detailed information about the installation, including location, tilt, azimuth, and maximum capacity. With this type of information, one can calculate the expected power output at each site and simply add up the sites [74]. Unfortunately, that information is seldom available and forecasters must take a top-down approach. Using this method, the AI irradiance forecast is trained for a site with available observations and then distributed geographically according to the distribution of the DPV [75-77]. Although this approach appears ad hoc, it has been found to produce forecasts within 3 to 5 percent of the actual at those sites where observations allow validation [77].

AI in the Longer Ranges

In the timeframe beyond one week, methods must be matched to the use. For the second week of the forecast, providers use methods quite similar to those discussed for the day-ahead timeframe—they work from NWP and correct with AI or statistical learning approaches. The subseasonal (>two weeks) to seasonal (one to three months) range begins to depend more on identifying the primary patterns in the atmosphere, often called teleconnection patterns, and predicting from observations of “typical” behavior during the prevailing events. Patterns such as the North American Oscillation and Pacific North American Oscillation denote “normal modes” of the atmosphere that can persist for a few weeks to a few months. At longer time scales, the lower temporal modes such as the El Niño/Southern Oscillation (ENSO) and Madden Julian Oscillation come into play. The problem transitions from one of an initial value problem solved by NWP models to more of a boundary value problem that must be solved via coupled systems that include more detailed interaction of the atmosphere with the oceans, land surface, cryosphere, and biosphere. National centers such as the U.S. National Oceanographic and Atmospheric Administration, Environment Canada, and European Center for Medium Range Forecasting routinely run seasonal forecasts using such coupled models. Forecasters often then use AI and statistical learning methods to improve on those forecasts.

Seasonal forecasts of renewable energy resources may be helpful to the military for long-term mission planning and base management as well as the hedging of financial risks through advanced purchasing of fuel. Gritti et al. [78] conducted experiments evaluating the accuracy of different learning methods for forecasting

whether the monthly wind and solar power generated at a geographically diverse set of sites across the United States would produce above or below normal power amounts. They compiled a large number of climate and teleconnection indices as input into a random forest model and compared the random forest with operational forecasts from the U.S. Climate Prediction Center. At most sites, the random forest predictions produced major improvements in accuracy over the simple benchmark and a random coin flip for both wind and solar anomalies. Utilizing a larger set of climate indices and predictions from seasonal climate models led to higher accuracy.

In the timeframes beyond a week, hydrological sources of renewable energy also vary substantially and forecasts help the energy industry plan their use. Those resources are typically modeled by coupling hydrological models with atmospheric models, such as the WRF-Hydro model (version of the WRF model adapted to interact with river segments) that was recently deployed as the U.S. National Water Model [79]. The hydrological community uses AI for reservoir management, model parameter estimation, snowfall calibration [80], and beyond.

For very long-range forecasts at the seasonal and longer timescales, the climate models are required, including all interactions discussed above, but additionally estimating changes in the atmospheric composition. To make estimates of interannual variability and projected changes of renewable resources, AI methods can facilitate a consistent comparison with current climate. For instance, in a collaborative project with the National Renewable Energy Laboratory, NCAR computed the interannual variability of the renewable resources as well as the projected change [81]. That study estimated the current climate variability using a climate reanalysis, which blends historical weather observations using an NWP model to intelligently interpolate the observations and ensure that the final analysis is consistent with model physics.

The team then used Self Organizing Maps (SOMs) to determine the primary modes of oscillation in the current climate. The future climate analysis was based on the output of a Global Climate Model (GCM) downscaled using an NWP model (called a Regional Climate Model [RCM] in this context). However, because the GCMs are known to lack some of the important teleconnection patterns such as ENSO, the RCM results were projected onto the SOMs of the current climate reanalysis. Then, a proxy climatology was built using a Monte Carlo simulation in which days from the current climate reanalysis were selected at the frequency modeled using the future climate RCM frequency for the same season and time of day. This proxy climate was then adjusted by appropriate “climate change” factors that were discerned in the difference between the reanalysis and the RCM model. In this way, the team showed that one can estimate future potential changes in the resource and model regional changes, but those projected changes tend to be within about 5 to 10 percent of the current resource [81]. Similar work has been accomplished in other regions of the world with similar resulting caps on the projected change.

Quantifying Uncertainty

Because the atmosphere is a nonlinear dissipative system, it is sensitive to initial conditions and inherently chaotic [82]. As a result, predictability is limited and meteorologists work to quantify the uncertainty in their forecasts, including for wind

speed and solar irradiance. The traditional approach to quantify that uncertainty is to take an ensemble forecasting approach. To generate an ensemble of potential solutions, modelers initialize multiple simulations by varying initial conditions, boundary conditions, physics parameterizations, or the models themselves [83]. The goal is to represent a reasonable probability density function (pdf) of the initial conditions that will evolve in time to a broader pdf representative of the forecasted flow. The statistics of that predicted pdf will then predict the uncertainty in the forecast. One typically finds that the mean of the ensemble produces a better forecast on average than any single member. Additionally, the ensemble provides probabilistic information that can be used to better integrate the renewable energy into the grid. For instance, some utilities are finding that with probabilistic predictions, they can dynamically allocate reserves more efficiently and economically. They can also predict potential distribution line congestion and thus take action to avoid frequency disruptions.

To fully represent the atmosphere would require hundreds of model simulations [84,85]. However, the NWP simulations are quite computationally expensive, so considerable research has been done to reduce that number using statistical learning and AI techniques. As the number becomes less than statistically ideal, the ensemble requires calibration. A pdf should be centered (the bias is removed), sharp (it distinguishes events), and reliable (the percentile predicted equals the percentile observed on average). Calibration acts to correct an imperfect ensemble to statistically represent the actual predicted variation. A straightforward method is Linear Variance Calibration [84,86]. This method recognizes that a reliable ensemble has an error variance that equals the variance of the ensemble pdf. So using historical data, the slope and intercept of the line fit between the two quantifies the error. Once those parameters are estimated, they can be used to correct, or calibrate, the ensemble in the future. Additional AI and statistical learning methods, including Bayesian Model Averaging [87,88] and quantile regression [58,89], among others, have also shown success at calibrating the uncertainty estimates.

Lee et al. [90,91] showed that if an NWP physics ensemble is clustered so as to distinguish the connections between members, one can eliminate redundancy by running merely a single member of each cluster and then calibrating the reduced ensemble. The research demonstrated that a 10-member ensemble that was calibrated using Bayesian Model Averaging produced comparable results to a full 42-member ensemble that was similarly calibrated.

Given that AI is useful for reducing the size of an ensemble, the next logical step is to design methods intelligently to forecast uncertainty with a single member. That goal has been accomplished using a method known as the analog ensemble (AnEn) [92,93]. The AnEn method uses a single high-quality NWP simulation with a sufficiently long time series (preferably a year or more) of historical predictions and their matching observations. The AnEn begins with today's forecast and searches the historical database for similar historical forecasts. Next, it identifies the validating observations for those most analogous forecasts and assembles them as a pdf, creating the AnEn. That ensemble is then used to both improve on the deterministic forecast and to quantify the uncertainty of the forecast. This method has been shown to produce an ensemble that is as correct and reliable as many of the best NWP ensembles of national centers [92,93], and has been

applied to wind energy [92,94] and solar power [95]. Figure 5 shows an example of three days of an AnEn solar irradiance forecast for a site near Sacramento, California. Not only does the AnEn improve upon the deterministic forecast, but it also allows the display of useful uncertainty information.




Figure 5. Example of three days forecast with the AnEn (yellow dashed) applied on top of the DICAST prediction (red) as compared to observations (black). Uncertainty bars are in blue.

A Look to the Future

With Advances in AI, Where Will This Go?

AI methods are certainly not new, but they have begun advancing again rapidly in the past 5–10 years, showing promise for all types of applications. These advances are starting to impact renewable energy applications, particularly forecasting. Some of the newer methods, such as the gradient-boosting methods, have rapidly become more mainstream and are being used more widely. Some of the newest AI advances include deep learning, which is beginning to be applied in a few places as discussed below.

The community is communicating regularly at conferences, workshops, and special sessions, which helps advance the state of the science. Some of the best methods blend knowledge of physics and dynamics to improve forecasts. This combination extends beyond using AI to optimize blending of physics models, as described above. It also includes intelligently choosing variables to use in training the AI models as well as how to derive variables based on physical relationships that may improve the output of the AI model.

The needs of the users are also becoming more sophisticated. Users have become used to relying on renewable forecast information and are starting to ask for more estimates of uncertainty in the forecast and variability forecasts to fine-tune their decisions. The meteorological and AI community are glad to comply as discussed below. We will see these AI applications enable yet more renewable energy.

Deep Learning and Its Potential Renewable Energy Forecasting Applications

Deep learning techniques have garnered a large amount of interest from the broader machine-learning community in recent years because of their success in image classification and speech translation [96]. Deep learning improves predictive accuracy by learning data representations with multiple levels of abstraction that discriminate between different classifications more distinctly [96]. Unlike traditional feedforward neural networks, deep-learning models, such as convolutional neural networks, contain multiple specialized layers that can identify features at different scales and levels of complexity and combine those features in an optimized fashion to produce more accurate predictions. Deep-learning methods provide the greatest amount of improvement over traditional machine learning methods when they are trained on minimally processed image or time series data because they build their own feature sets from the raw data. Convolutional neural networks, which learn sets of convolutional feature maps applied to gridded data to transform that data into a more interpretable format, have shown a lot of promise for analyzing the observational and gridded forecast datasets used in meteorology and other earth sciences [97].

Convolutional neural networks have the potential to improve predictions at all forecast lead times for renewable energy predictions. For very short-term predictions, convolutional neural networks can identify cloud types from sky camera data [98]. Training a convolutional neural network from scratch requires a very large image dataset, which is often not available for most specialized problems. However, many images can be described well with similar sets of small-scale features, so researchers have transferred the convolutional layers from networks trained on a general set of images and only trained the fully connected layers on the smaller set of specialized data. In Ye et al. [98], this approach yielded large gains in accuracy compared with existing cloud identification methods.

Convolutional neural networks could also be used to identify spatial features in gridded meteorological data and translate those features into renewable energy predictions. By identifying the presence of weather patterns favorable for cloudiness or rapid changes in wind speed, the convolutional neural network may be able to generate more informed predictions than models relying on model output from the nearest grid point to a location of interest. In particular, convolutional neural networks could identify the combinations of fine- and coarse-scale weather features that result in rapid changes in wind and solar irradiance.

Forecasting Variability, Quantifying Uncertainty

Most AI renewable energy forecasting systems focus on the average value of wind speed or solar irradiance over the time periods of interest, but the average value does not provide users with information about the number of large fluctuations possible within a given period. If large fluctuations in power are more likely to occur, then the electricity provider would require more reserve power to balance the load [47]. Because power variability also depends on the local weather conditions at the generation site, similar AI techniques and input variables can be used to forecast variability as are currently used to forecast average power. The challenge is determining the most appropriate variability metric for a user. McCandless et al. [47] predicted the 15-minute spatial and temporal standard deviations of solar

irradiance using a decision-tree regression model and produced consistent error amounts in the one to three-hour lead time range. Lave et al. [99] quantified solar irradiance variability in terms of the distribution of ramp rates, which describe how much the power or irradiance changes between measurements. Sites in Hawaii and Puerto Rico recorded more variability on average than sites in California and Arizona [99]. Standard deviation is easier to interpret, but it does not provide any information about the frequency of changes, whereas the ramp-rate distribution accounts for those temporal changes.

Variability forecasts should not be confused with estimates of forecast uncertainty. The amount of uncertainty is a function of the forecast system and its design, while the amount of variability is a function of the observed phenomena and how it is observed. Uncertainty is dependent on the lead time of the forecast, the amount of available information, and the configuration of the forecast model being used. The amount of variability is dependent only on the space and time intervals specified for aggregation. Variability forecasts can be deterministic, or they can incorporate additional uncertainty information. Uncertainty estimates around forecasts of a mean quantity are not guaranteed to reflect the variability of that quantity within a given time period. End users should carefully consider what uncertainty and variability information they need from their power forecasts to be prepared for fluctuations in their electrical generation system.

As storage becomes more available, it will allow smoothing of the variability of the renewable resources. However, to optimize when to store and when to dispatch power will still require accurate forecasts.

Summary and Conclusions

Because renewable energy is growing rapidly, it is impacting operations of the energy sector as well as DoD's planning and use of power. Renewable resources are variable and depend on weather information; thus, its efficient and economic use depends critically on having accurate forecasts of available power at multiple time scales to enable using it to balance electrical load. At the short scale of the next several hours, forecasts are required for grid integration and ensuring a reliable and economic energy supply. For day-ahead unit allocation and planning, one must have forecasts on these scales. Longer-term planning for maintenance and deployment requires forecasts on scales of weeks, months, seasons, and decades. At all of these scales, optimal predictions rely on knowledge of the physics plus skilled use of AI and statistical learning methods to optimize the forecasts.

The meteorological and computing communities have collaborated to produce methods to optimize forecasts, using forecasting technology from statistics, AI, and NWP. The AI methods are advancing prediction capabilities beyond the accuracy of the physics alone.

But atmosphere flow is inherently chaotic, implying a limit to predictability. Thus, as the skill of the forecast improves, operators are learning to use information on the variability of the resource and uncertainty in the forecasts, while also finding other uses for such information. As technology increases in complexity, so do its applications.

References

1. Renewable Energy Policy Network for the 21st Century. (2017). *Global status report highlights: Advancing the global renewable energy transition*. Retrieved from http://www.ren21.net/wp-content/uploads/2017/06/GSR2017_Highlights_FINAL.pdf
2. Rapier, R. (2016, June 3). A record year for renewable energy. *Forbes*. Retrieved from <https://www.forbes.com/sites/rrapier/2016/06/03/a-record-year-for-renewable-energy/#3285707b6c90>
3. American Wind Energy Association. (2017). Wind energy facts at a glance. Retrieved from <http://www.awea.org/wind-energy-facts-at-a-glance>
4. Solar Energy Industries Association. (n.d.). Solar industry data. Retrieved from <http://www.seia.org/research-resources/solar-industry-data>
5. International Energy Agency (n.d.). Renewables statistics. Retrieved from <http://www.iea.org/statistics/topics/renewables>
6. Dubus, L. (2014). Weather and climate and the power sector: Needs, recent developments, and challenges. In A. Troccoli, L. Dubus, & S.E. Dubus (Eds.), *Weather matters for energy* (pp. 379–398). New York, NY: Springer. doi:10.1007/978-1-4614-9221-4_18
7. Lew, D., Brinkman, G., Florita, A., Heaney, M., Hodge, B-M., Hummon, M., & Ibanez, E. (2012, November). *Sub-hourly impacts of high solar penetrations in the western United States*. Presented at the 2nd Annual International Workshop on Integration of Solar Power into Power Systems Conference, Lisbon, Portugal.
8. Ela, E., Diakov, V., Ibanez, E., & Heaney, M. (2013). *Impacts of variability and uncertainty in solar photovoltaic generation at multiple timescales* (NREL/TP-5500-58274). Golden, CO: National Renewable Energy Laboratory. Retrieved from <http://www.nrel.gov/docs/fy13osti/58274.pdf>
9. Curtright, A.E., Apt, J. (2008). The character of power output from utility-scale photovoltaic systems. *Progress in Photovoltaics: Research and Applications*, 16(3), 241–247. doi:10.1002/pip.786
10. Kosovic, B., Haupt, S.E., Adriaansen, D., Alessandrini, S., Jensen, T., Wiener, G.,... Delle Monache, L. (n.d.). Scientific advances in wind power forecasting. Manuscript submitted for publication.
11. Ahlstrom, M., Bartlett, D., Collier, C., Duchesne, J., Edelson, D., Gesino, A., . . . Rodriguez M. (2013). Knowledge is power: Efficiently integrating wind energy and wind forecasts. *IEEE Power and Energy Magazine*, 11(6), 45–52. doi:10.1109/mpe.2013.2277999
12. Tuohy, A., Zack, J., Haupt, S. E., Sharp, J., Ahlstrom, M., Dise, S., . . . Collier, C. (2015). Solar forecasting: Methods, challenges, and performance. *IEEE Power and Energy Magazine*, 13(6), 50–59. doi:10.1109/MPE.2015.2461351
13. Orwig, K. D., Ahlstrom, M. L., Banunarayanan, V., Sharp, J., Wilczak, J. M., Freedman, J., . . . Marquis, M. (2014). Recent trends in variable generation forecasting and its value to the power system. *IEEE Transactions on Sustainable Energy*, 6(3), 924–933. doi:10.1109/TSTE.2014.2366118

14. U.S. Department of Defense. (2014, April 16). *Department of Defense Directive Number 4180.01 - DoD Energy Policy*. Retrieved from https://fas.org/irp/doddir/dod/d4180_01.pdf
15. Environmental and Energy Study Institute. (2011, July 27). Fact sheet - DoD's energy efficiency and renewable energy initiatives. Retrieved from <http://www.eesi.org/papers/view/fact-sheet-dods-energy-efficiency-and-renewable-energy-initiatives>
16. U.S. Department of Defense. (2016). *Department of Defense 2016 operational energy strategy*. Retrieved from http://www.acq.osd.mil/eie/Downloads/OE/2016%20OE%20Strategy_WEBd.pdf
17. Vaucher, G. (2015, September). *Atmospheric renewable-energy research, Volume 1 (Background: "To BE or not to BE")* (ARL-TR-7402). White Sands Missile Range, NM: U.S. Army Research Laboratory. Retrieved from <http://www.dtic.mil/docs/citations/ADA622364>
18. Vaucher, G. (2016). *Atmospheric renewable-energy research, Volume 2: Assessment process for solar powered meteorological applications* (ARL-TR-7762). White Sands Missile Range, NM: U.S. Army Research Laboratory.
19. Saritas, O., & Burmaoglu, S. (2016). Future of sustainable military operations under emerging energy and security considerations. *Technological Forecasting and Social Change, 102*, 331–343. doi:10.1016/j.techfore.2015.08.010
20. Eady, D. S., Siegel, S. B., Bell, R. S., & Dicke, S. H. (2009, September). *Final technical report: sustain the mission project: Casualty factors for fuel and water resupply convoys* (CTC-CR-2009-163). Arlington, VA: Army Environmental Policy Institute. Retrieved from <http://www.dtic.mil/dtic/tr/fulltext/u2/b356341.pdf>
21. CNA Military Advisory Board. (2017). *Advanced energy and U.S. national security*. Retrieved from https://www.cna.org/CNA_files/PDF/IRM-2017-U-015512.pdf
22. Kleissl, J. (2013). *Solar energy forecast and resource assessment* (1st ed.). Amsterdam: Academic Press.
23. Troccoli, A., Dubus, L., & Haupt, S.E. (Eds.). (2014). *Weather matters for energy*. New York, NY: Springer.
24. Mellit, A. (2008). Artificial intelligence technique for modeling and forecasting of solar radiation data: A review. *International Journal of Artificial Intelligence and Soft Computing, 1*(1), 52–76. doi:10.1504/IJAISC.2008.021264
25. Martín, L., Zorzalejo, L. F., Polo, J., Navarro, A., Marchante, R., & Cony, M. (2010). Prediction of global solar irradiance based on time series analysis: Application to solar thermal power plants energy production planning. *Solar Energy, 84*(10), 1772–1781. doi:10.1016/j.solener.2010.07.002
26. Pedro, H. T. C., & Coimbra, C. F. M. (2012). Assessment of forecasting techniques for solar power prediction with no exogenous inputs. *Solar Energy, 86*, 2017–2028. doi:10.1016/j.solener.2012.04.004
27. Notton, G., Paoli, C., Vasileva, S., Nivet, M-L., Canaletti, J-L., & Cristofari, C. (2012). Estimation of hourly global solar irradiation on

- tilted planes from horizontal one using artificial neural networks. *Energy*, 39(1), 166–179. doi:10.1016/j.energy.2012.01.038
28. Bhardwaj, S., Sharma, V., Srivastava, S., Sastry, O. S., Bandyopadhyay, B., Chandel, S., & Gupta, J. R. P. (2013). Estimation of solar radiation using a combination of Hidden Markov Model and generalized Fuzzy model. *Solar Energy*, 93, 43–54. doi:10.1016/j.solener.2013.03.020
 29. Bouzerdoum, M., Mellit, A., & Pavan, A. M. (2013). A hybrid model (SARIMA-SVM) for short-term power forecasting of a small-scale grid-connected photovoltaic plant. *Solar Energy*, 98, 226–235. doi:10.1016/j.solener.2013.10.002
 30. Diagne, M., David, M., Lauret, P., Boland, J., & Schmutz, N. (2013). Review of solar irradiance forecasting methods and a proposition for small-scale insular grids. *Renewable and Sustainable Energy Reviews*, 27, 65–76. doi:10.1016/j.rser.2013.06.042
 31. Fu, C-L., & Cheng, H-Y. (2013). Predicting solar irradiance with all-sky image features via regression. *Solar Energy*, 97, 537–550. doi:10.1016/j.solener.2013.09.016
 32. Marquez, R., Pedro, H.T.C., & Coimbra, C.F.M. (2013). Hybrid solar forecasting method uses satellite imaging and ground telemetry as inputs to ANNs. *Solar Energy*, 92, 176–188. doi:10.1016/j.solener.2013.02.023
 33. Inman, R. H., Pedro, H. T. C., & Coimbra, C. F. M. (2013). Solar forecasting methods for renewable energy integration. *Progress in Energy and Combustion Science*, 39(6), 535–576. doi:10.1016/j.pecs.2013.06.002
 34. Chu, Y., Pedro, H. T., & Coimbra, C. F. M. (2013). Hybrid intra-hour DNI forecasts with sky image processing enhanced by stochastic learning. *Solar Energy*, 98, 592–603. doi:10.1016/j.solener.2013.10.020
 35. Fernandez, E. F., Almonacid, F., Sarmah, N., Rodrigo, P., Mallick, T. K., & Pérez-Higueras, P. (2014). A model based on artificial neuronal network for the prediction of the maximum power of a low concentration photovoltaic module for building integration. *Solar Energy*, 100, 148–158. doi:10.1016/j.solener.2013.11.036
 36. Almonacid, F., Pérez-Higueras, P. J., Fernández, E. F., & Hontoria, L. (2014). A methodology based on dynamic artificial neural network for short-term forecasting of the power output of a PV generator. *Energy Conversion and Management*, 85, 389–398. doi:10.1016/j.enconman.2014.05.090
 37. Quesada-Ruiz, S., Chu, Y., Tovar-Pescador, J., Pedro, H. T. C., & Coimbra, C. F. M. (2014). Cloud-tracking methodology for intra-hour DNI forecasting. *Solar Energy*, 102, 267–275. doi:10.1016/j.solener.2014.01.030
 38. Reikard, G., Haupt, S. E., & Jensen, T. (2017). Forecasting ground-level irradiance over short horizons: Time series, meteorological, and time-varying parameter models. *Renewable Energy*, 112, 474–485. doi: 10.1016/j.renene.2017.05.019
 39. López, G., Battles, F.J., & Tovar-Pescador, J. (2005). Selection of input parameters to model direct solar irradiance by using artificial

- neural networks. *Energy*, 30(9), 1675–1684.
doi:10.1016/j.energy.2004.04.035
40. Greybush, S. J., Haupt, S. E., Young, G. S. (2008). The regime dependence of optimally weighted ensemble model consensus forecasts of surface temperature. *Weather and Forecasting*, 23(6), 1146–1161. doi:10.1175/2008waf2007078.1
 41. Hinkelman, L.M. (2014). Relating solar resource variability to cloud type. NREL Technical Report.
 42. Tapakis, R., & Charalambides, A.G. (2013). Equipment and methodologies for cloud detection and classification: A review. *Solar Energy*, 95, 392–430.
 43. McCandless, T.C., Haupt, S. E., & Young, G.S. (2014). Short term solar radiation forecasts using weather regime-dependent artificial intelligence techniques. 12th Conference on Artificial Intelligence: Applications of Artificial Intelligence Methods for Energy. Atlanta, GA. Retrieved from <https://ams.confex.com/ams/94Annual/webprogram/Paper240879.html>
 44. Morf, H. (2014). Sunshine and cloud cover prediction based on Markov processes. *Solar Energy*, 110, 615–626.
 45. Zagouras, A., Kazantzidis, A., Nikitidou, E., & Argiriou, A.A. (2013). Determination of measuring sites for solar irradiance, based on cluster analysis of satellite-derived cloud estimations. *Solar Energy*, 97, 1–11.
 46. Mellit, A., Massi Pavan, A., & Lughi, V. (2014). Short-term forecasting of power production in a large-scale photovoltaic plant. *Solar Energy*, 105, 401–413.
 47. McCandless, T.C., Haupt, S. E., & Young, G.S. (2015). A model tree approach to forecasting solar irradiance variability. *Solar Energy*, 120, 514–524. doi:10.1016/j.solener.2015.07.020
 48. Chow, C.W., Urquhart, N., Lave, M., Dominquez, A., Kleissl, J., Shields, J., & Washom, B. (2011). Intra-hour forecasting with a total sky imager at the UC San Diego solar energy testbed. *Solar Energy*, 85, 2881–2893.
 49. Marquez, R., & Coimbra, C.F.M. (2013a). Intra-hour DNI forecasting based on cloud tracking image analysis. *Solar Energy*, 91, 327–336.
 50. Huang, H., Xu, J., Peng, Z., Yoo, S., Yu, D., Huang, D., & Qin, H. (2013, October 21–24). Cloud motion estimation for short term solar irradiance prediction. 2013 IEEE International Conference on Smart Grid Communications.
 51. Chu, Y., Pedro, H.T.C., Li, M., & C.F.M. Coimbra (2015). Real-time forecasting of solar irradiance ramps with smart image processing. *Solar Energy*, 114, 91–104.
 52. Peng, Z., Yu, D., Huang, D., Heiser, J., Yoo, S., & Kalb P. (2015). 3D cloud detection and tracking system for solar forecast using multiple sky imagers. *Solar Energy*, 118, 496–519, doi:10.1016/j.solener.2015.05.037
 53. Lorenz, E., Hammer, A., & Heinemann, D. (2004). Short term forecasting of solar radiation based on satellite data. EUROSUN2004, ISES Europe Solar Congress, 841–848.
 54. Miller, S.D., Forsythe, J.M., Partain, P.T., Haynes, J.M., Bankert, R.L., Sengupta, M., ... & Vonder Haar, T.H. (2014). Estimating three-

- dimensional cloud structure via statistically blended satellite observations. *Journal of Applied Meteorology and Climatology*, 53, 437–455.
55. Cros, S., Liandrat, O., Sébastien, N., & Schmutz, N. (2014). Extracting cloud motion vectors from satellite images for solar power forecasting. *2014 IEEE Geoscience and Remote Sensing Symposium*. doi:10.1109/igarss.2014.6947394
 56. McCandless, T.C., Haupt, S.E., & Young G.S. (2016a). A regime-dependent artificial neural network technique for short-range solar irradiance forecasting. *Renewable Energy*, 89, 351–359. doi:10.1016/j.renene.2015.12.030
 57. McCandless, T.C., Young, G.S., Haupt, S.E., & Hinkelman, L.M (2016b). Regime-dependent short-range solar irradiance forecasting. *Journal of Applied Meteorology and Climatology*, 55, 1599–1613.
 58. Mahoney, W. P., Parks, K., Wiener, G., Liu, Y., Myers, W. L., Sun, J., . . . Haupt, S. E. (2012). A wind power forecasting system to optimize grid integration. *IEEE Transactions on Sustainable Energy*, 3(4), 670–682. doi:10.1109/tste.2012.2201758
 59. Voyant, C., Muselli, M., Paoli, C., & Nivet, M-L. (2014). Numerical weather prediction (NWP) and hybrid ARMA/ANN to predict global radiation. *Energy*, 39(1), 341–355.
 60. Voyant, C., Muselli, M., Paoli, C., & Nivet, M-L. (2013). Hybrid methodology for hourly global radiation forecasting in Mediterranean area. *Renewable Energy*, 53, 1–11.
 61. Haupt, S.E. & Kosovic, B. (2016). Variable generation power forecasting as a big data problem. *IEEE Transactions on Sustainable Energy*, 8(2), 725–732. doi:10.1109/TSTE.2016.2604679
 62. Haupt, S.E., Jimenez, P., Lee, J.A., & Kosovic, B. (2017c). *Principles of meteorology and numerical weather prediction, in renewable energy forecasting: From models to applications*. London: Elsevier. In press.
 63. Jimenez, P.A., Hacker, J.P., Dudhia, J., Haupt, S.E., Ruiz-Arias, J.A., Gueymard, C.A., . . . & Deng, A.J. (2016a, July). WRF-Solar: An augmented NWP model for solar power prediction. *Bulletin of the American Meteorological Society*, 1249–1264. doi:10.1175/BAMS-D-14-00279.1
 64. Jimenez, P.A., Alessandrini, S., Haupt, S.E., Deng, A., Kosovic, B., Lee, J.A., & Delle Monache, L. (2016b). Role of unresolved clouds on short-range global horizontal irradiance predictability. *Monthly Weather Review*, 144, 3099–3107. doi:10.1175/MWR-D-16-0104.1
 65. Fitch, A. C., Olson, J. B., Lundquist, J. K., Dudhia, J., Gupta, A. K., Michalakes, J., & Barstad, I. (2012). Local and mesoscale impacts of wind farms as parameterized in a mesoscale NWP model. *Monthly Weather Review*, 140(9), 3017–3038. doi:10.1175/mwr-d-11-00352.1
 66. Glahn, H. R., & Lowry, D. A. (1972). The use of model output statistics (MOS) in objective weather forecasting. *Journal of Applied Meteorology*, 11, 1203–1211. doi:10.1175/1520-0450(1972)011<1203:TUOMOS>2.0.CO;2
 67. Wilks, D.S. (2006). *Statistical methods in the atmospheric sciences*, 2nd Ed., Academic Press, Burlington, MA.

68. Myers, W., Wiener, G., Linden, S., and Haupt, S.E. (2011). A consensus forecasting approach for improved turbine hub height wind speed predictions. *Proceedings of WindPower 2011*, Anaheim, CA.
69. Haupt, S.E. & Mahoney, W.P. (2015, November). Wind power forecasting. *IEEE Spectrum*, 46–52.
70. Haupt, S. E., Kosović, B., Jensen, T., Lazo, J. K., Lee, J. A., Jiménez, P. A., . . . Heiser, J. (2018). Building the Sun4Cast System: Improvements in solar power forecasting. *Bulletin of the American Meteorological Society*, 99(1), 121-136. doi:10.1175/bams-d-16-0221.1
71. Haupt, S. E., & Coauthors (2016a). The Sun4Cast solar power forecasting system: The results of the public-private-academic partnership to advance solar power forecasting. pp. doi:10.5065/D6N58JR2
72. Gagne, D.J., McGovern, A., Haupt, S.E., & Williams, J.K. (2017). Evaluation of statistical learning configurations for gridded solar irradiance forecasting. *Solar Energy*, 150, 383–393. doi:10.1016/j.solener.2017.04.031
73. McGovern, A., Gagne, D.J., Basara, J., Hamill, T.M., & Margolin, D. (2015). Solar energy prediction: An international contest to initiate interdisciplinary research on compelling meteorological problems. *Bulletin of the American Meteorological Society*, 96, 1388–1395. doi:10.1175/BAMS-D-14-00006.1
74. Hoff, T. (2016, September 28). Behind-the-meter (BTM) PV Forecast Integration with Load Forecasting, Utility Variable Generation Integration Group Forecasting Workshop. Denver, CO.
75. Lorenz, E., Kuhnert, J., & Heinemann, D. (2012). Overview on irradiance and photovoltaic power prediction. In A. Troccoli, L. Dubus, & S.E. Dubus (Eds.), *Weather Matters for Energy*, 429–454. New York: Springer. doi:10.1007/978-1-4614-9221-4_18
76. Gagne, D.J., Haupt, S.E., Linden, S., Williams, J.K., McGovern, A., Wiener, G., . . . & McCandless, T.C. (2015). Scaling machine learning models to produce high resolution gridded solar power forecasts. 13th Conference on Artificial Intelligence: The Last Mile: Methods and Technologies for Delivering Custom Weather, Water, and Climate Information to Everyone in the World. Phoenix, AZ.
77. Haupt, S. E., Dettling, S., Williams, J. K., Pearson, J., Jensen, T., Brummet, T., . . . Burghardt, C. (2017). Blending distributed photovoltaic and demand load forecasts. *Solar Energy*, 157, 542-551. doi:10.1016/j.solener.2017.08.049
78. Gritmit, E.P., Sinilo, A., Stoelinga, M.T., McCaa, J., and Storck, P. (2017). Benchmarking the accuracy of seasonal forecasting for renewable energy resource anomalies. 8th Conference on Weather, Climate, Water and the New Energy Economy. Seattle, WA. Retrieved from <https://ams.confex.com/ams/97Annual/webprogram/Paper314192.htm>
79. Gochis, D. J., Yu, W., & Yates, D.N. (2015). The WRF-Hydro modeling system Retrieved from https://www.ral.ucar.edu/projects/wrf_hydro

80. McCandless, T.C., Haupt, S. E., & Young, G.S. (2011). Statistical guidance methods for predicting snowfall accumulation in the northeast United States. *National Weather Digest*, 35, 149–162.
81. Haupt, S.E., Copeland, J., Cheng, W.Y.Y., Amman, C., Zhang, Y., & Sullivan, P. (2016b). Quantifying the wind and solar power resource and their inter-annual variability over the United States under current and projected future climate. *Journal of Applied Meteorology and Climatology*, 55, 345–363. doi:10.1175/JAMC-D-15-0011.1
82. Lorenz, E. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20, 130–141.
83. Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. NY: Cambridge University Press.
84. Kolczynski, W.C., Stauffer, D.R., Haupt, S.E., & Deng, A. (2009). Ensemble variance calibration for representing meteorological uncertainty for atmospheric transport and dispersion modeling. *Journal of Applied Meteorology and Climatology*, 48(10), 2001–2021. doi:10.1175/2009jamc2059.1
85. Wang, X., and Bishop, C.H. (2005). Improvement of ensemble reliability with a new dressing kernel. *Quarterly Journal of the Royal Meteorological Society*, 131, 965–986.
86. Kolczynski, W.C., Stauffer, D.R., Haupt, S.E., Altman, N.S., & Deng, A. (2012). Investigation of linear variance calibration for spread-error relationship using a stochastic model. *Monthly Weather Review*, 139, 3954–3963.
87. Raftery, A.E, Balabdaoui, F., Gneiting T., & Polakowski, M. (2003, December 15). *Using Bayesian model averaging to calibrate forecast ensembles* (Technical Rep. No. 40). Department of Statistics, University of Washington. Retrieved from <http://www.stat.washington.edu/research/reports/2003/tr440.pdf>
88. Raftery, A.E, Balabdaoui, F. Gneiting, T., & Polakowski, M. (2005). Using Bayesian model averaging to calibrate forecast ensembles.
89. Hopson, T.M. (2014). Assessing the ensemble spread-error relationship, *Monthly Weather Review*, 142, 1125–1142. doi:10.1175/MWR-D-12-00111.1
90. Lee, J.A., Haupt, S.E., & Young, G.S. (2016). Down-selecting NWP multi-physics ensembles with hierarchical cluster analysis. *Journal of Climatology and Weather Forecasting*, 4(1). doi:10.4172/2332-2594.1000156
91. Lee, J.A., Kolczynski, W.C., McCandless, T.C., and Haupt, S.E. (2012). Objective techniques for configuring and down-selecting an NWP ensemble for low-level wind predictions. *Monthly Weather Review*, 140, 2270–2286.
92. Delle Monache, L., Nipen, T., Liu, Y., Roux, G., & Stull, R. (2011). Kalman filter and analog schemes to postprocess numerical weather predictions. *Monthly Weather Review* 139, 3554–3570. doi:10.1175/2011MWR3653.1
93. Delle Monache, L., Eckel, F.A., Rife, D.L., Nagarajan, B., & Searight, K. (2013). Probabilistic weather prediction with an analog ensemble. *Monthly Weather Review*, 141, 3498–3516. doi:10.1175/MWR-D-12-00281.1

94. Haupt, S.E. & Delle Monache, L. (2014). Understanding ensemble prediction: How probabilistic wind power prediction can help in optimizing operations. *WindTech*, 10, 27–29.
95. Alessandrini, S., Delle Monache, L., Sperati, S., & Cervone, G. (2015). An analog ensemble for short-term probabilistic solar power forecast. *Applied Energy*, 157, 95–110, doi:10.1016/j.apenergy.2015.08.011
96. LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444, doi:10.1038/nature14539
97. Dieleman, S., Willett, K.W., & Dambre, J. (2015). Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Monthly Notices of the Royal Astronomical Society*, 450(4), 1441–1459. doi:10.1093/mnras/stv632.
98. Ye, L., Cao, Z., & Xiao, Y. (2017, June 30). DeepCloud: Ground-based cloud image categorization using deep convolutional features. *IEEE Transactions on Geoscience and Remote Sensing*, 1–12, doi:10.1109/TGRS.2017.2712809
99. Lave, M., Reno, M.J., & Broderick, R.J. (2015). Characterizing local high-frequency solar variability and its impact to distribution studies. *Solar Energy*, 118, 327–337. doi:10.1016/j.solener.2015.05.028

3

Biometrics, Critical Infrastructure Protection, and Homeland Defense and Security

Anthony Rucci, Information International Associates

Gregory Nichols, Homeland Defense & Security Information Analysis Center

Applications of AI across biometrics, CIP, and HDS have the potential to integrate so well that it seems appropriate to discuss all three topics at once. For example, DHS is testing AVATAR, which uses AI algorithms to screen passengers at border crossings and airports [1]. AVATAR detects physiological responses and biometric features to determine if a passenger is trustworthy. Once a person passes this test, they can move on in their journey; if they do not pass, then a secondary screening is performed by a human. Utilizing technologies such as AVATAR to identify bad actors attempting to infiltrate airports and airplanes might help improve the safety of Transportation Systems, considered one of DHS's 16 critical infrastructure sectors [2]. The protection of critical infrastructure, such as commercial aviation, falls under the aegis of homeland security: "the national effort to ensure a homeland that is safe, secure, and resilient against terrorism and other hazards where American interests, aspirations, and ways of life can thrive to the national effort to prevent terrorist attacks within the United States, reduce the vulnerability of the U.S. to terrorism, and minimize the damage from attacks that do occur [3]."

Advances in AI have led to numerous tools that are useful to the DoD and military for homeland defense. AI's ability to recognize patterns in culture, conflict, and crises creates an avenue for progressive simulation and autonomous aid. Much of AI use for homeland defense comes in the form of simulating specific scenarios. Using AI could be useful in general avenues of research such as simulation for military planning, simulation for warfighter and troop performance enhancement, and programming for autonomous vehicles and robotics for deployment in areas unsafe for humans. AI could be instrumental in a wide area of military life—from running computer-based simulation strategy for war-games and assessing the mental health of soldiers, to modeling battlefield scenarios that gauge psychological functioning and decision-making in high intensity environments.

A lesser-known application for AI use in the DoD is computational social science programs that can aid in military operations [4]. The creation of training platforms enhances interpersonal cohesion such that soldiers from varied backgrounds are able to work well as a unit [4]. AI based simulation has been extended to characterizing the behavioral signatures of adversaries through modeling complex historical conflicts and current global conditions [5].

Recognition-primed decision (RPD) theory attempts to define the thought process that effective decision makers employ in complex situations [6]. The development of AI based pattern recognition algorithms coupled with tenants of RPD theory allow for aid in crisis management and risk assessment practices. Applications include HAZMAT emergency and firefighting responses [7] and the ability to predict challenges to national security by violent actors aiming to do harm through counterinsurgency and terrorism [8].

Implementation of AI in autonomous machines provides an avenue for advanced military robotics. Generally, military robots are defined as mechanical systems that (1) possess the capacity to sense their surroundings, communicate information, and make decisions; (2) can act autonomously according to predetermined rules or under human supervision; (3) are capable of improving their own performance through machine learning [9]. The addition of AI to these systems will provide military robotics with the ability to make real time decisions on a plan of action based on pattern recognition derived from historical and mission critical information caches. AI-supplemented military robots will support reconnaissance, explosives detonation, medical evacuation, and advanced weaponry [9].

A significant part of HDS involves weapons development. In terms of weaponeering, AI is currently used in over 280 weapons systems [10], with the majority of these applications being homing capabilities. Since 1965, autonomy and related technologies, including AI, in weapons platforms have increased in general. Weapons development is primarily covered by the Defense Systems Information Analysis Center, so the topic will only briefly be covered here; however, it is worth discussing some of the concerns of AI weaponry and related defense challenges, as these new technologies could pose a security risk to the U.S. homeland in the future. DoD is already exploring options to combat the use of AI-capable weapons. In December 2015, during a speech hosted by the Center for a New American Security, Deputy Secretary of Defense Bob Work commented on the Third Offset and stated, “Now, we believe, strongly, that humans should be the only ones to decide when to use lethal force. But when you’re under attack, especially at machine speeds, we want to have a machine that can protect us [11].”

In October 2016, at the Center for Strategic and International Studies, Secretary Work remarked: “The pacing competitors—not adversaries—are Russia and China, because they’re developing advanced capabilities that potentially worry us [12].” This statement appears to ring true as China and Russia are investing heavily in AI research for military applications. In August 2016, Wang Changqing, director of the General Design Department of the Third Academy of the China Aerospace Science and Industry Corporation, claimed Chinese engineers have spent years researching the use of AI in missiles [13], and they are currently developing intelligent cruise missiles similar to the Tomahawk. In the 1970s, Russia developed one of the first missiles to use AI, the Granit missile (SS-N-19 Shipwreck). They are now replacing the Granit with smaller AI-guided P-800 Oniks missiles (SS-N-26 Strombile). The smaller size means almost three times as many missiles can be loaded on to ships [14].

In testimony to the Senate Armed Services Committee in 2016, Admiral Harry Harris, commander of USPACOM, requested that Congress authorize the purchase of the new Long Range Anti-Ship Missile (LRASM) [15]. The LRASM is completely autonomous, so after a human operator tells the missile which ship to target and where the fleet is, AI handles the rest and uses a continuous stream of data to correct course until it hits its target.

As an example of military robotics, in April 2017, Russia released video footage of their Final Experimental Demonstration Object Research (FEDOR), a bipedal robot capable of duplicating many human movements such as driving, giving first

aid, handling tools, and even shooting. Russia claims that FEDOR will replace astronauts on the International Space Station [16].

Efforts to counter AI weaponry are under development as the U.S. emphasizes machine learning. In addition, AI is employed in other platforms to combat non-AI weaponry. The Synchronized Net-Enabled Multi-INT Exploitation utilizes “advanced intelligence-collection and rapid-response capability” to identify weaknesses in an enemy’s operations [17]. In July 2017, Colonel Drew Cukor, head of the Algorithmic Warfare Cross Function Team, stated that the U.S. will be using algorithms by the end of 2017 to sort through intelligence data and make decisions for combating groups such as the Islamic State more efficiently and effectively [18].

AI is increasingly used to process biometric data. Humans have a number of highly individualized characteristics, most notably their fingerprints. Historically, latent print examiners interpreted fingerprints for identity verification, but a large-scale study conducted by Carnegie Mellon University reports that 85 percent of examiners made at least one false negative error [19]. This study exemplifies the need for advancing biometric identification technology. The research conducted on biometric analysis pinpoints a variety of personal characteristics that can greatly increase the accuracy of identification protocol.

The biometric identification process involves matching user input data—a fingerprint scan for example—with information from a specified cache of data. Algorithms will sort through the cache of data and search for information that matches what the user has input. Limitations of biometric technology are reliable data caches, processing power, and the resolution of input data. The pattern recognition and predictive capabilities of AI algorithms are well suited for information processing.

Global reliance on technology has increased the amount of available digital information. Sifting through this information rapidly is necessary for real time biometric identification. The DoD is addressing this problem by narrowing the cache of information that is necessary to process. Catalogs contain visual, audio, and hyperspectral data through which an AI algorithm can search. This functionality enables multimodal biometric identification by matching personal characteristics such as facial structure, height, gait, and skin pigmentation recognition [20]. Audio files containing examples of an individual’s voice and speech patterns can be used as biometric identification tools as well [21, 22]. Even information not typically thought to be personal characteristics such as keystroke pattern, hand geometry, vascular orientation, and clothing recognition are used as platforms for AI biometric identification [23-25].

AI is useful in feature selection, the act of discarding unnecessary data in order to reduce processing time and increase identification accuracy. For example, an audio file of a conversation between two people in a crowded restaurant contains background noise such as the clanking of silverware, people from surrounding tables, etc. Elimination of the irrelevant data is necessary for speaker identification, which is where AI excels. Machine learning allows AI algorithms to perform tasks with increasing accuracy after each iteration. In other words, machine learning

helps the AI to continually recognize and extract information on specified features with a higher degree of certainty.

Future applications of AI-assisted biometric recognition technologies might include passenger screening at airports, border security, and identification of persons of interest to the DoD [26]. A U.S. Air Force study ranked image processing and AI in their top 11 of 43 key technologies to develop by 2025 [27]. Extensions of biometric identification will enhance battlefield support through AI-supported combatant recognition and smart weaponry [26], as demonstrated by the Super aEgis II .50 caliber machine gun, which has the capability to recognize and target human faces [28].

Integrating AI into CIP, particularly cyber-physical systems, has been more challenging than the incorporation of AI into other areas, primarily due to safety and security concerns regarding the possibility of catastrophic failure. A recent DHS report [29] acknowledges the benefits that AI could play in critical infrastructure protection. The report also noted several risks from adoption and potential safety measures to mitigate those risks. Consider the context of using monitoring capabilities that already exist in traditional networks and adding AI technology to determine anomalous activity. Based on its telemetry, AI might take action autonomously based on a predetermined system setting, a range of acceptable norms sensing change, or a “perceived change” in the data being ingested into that controller. The indicators of compromise would trigger pre-set responses without human intervention. While on the surface this might sound like an exceptionally efficient model, a question remains: if unsupervised, what actions might the system take that might trigger a catastrophic chain of events?

The ultimate reward for an adversary would be to disrupt one or more segments of U.S. critical systems. Our enemies will likely place a lot of effort in compromising critical nodes such as our Industrial Control Systems, Supervisory Control and Data Acquisition, Distributed Control Systems, or Programmable Logic Controllers. Since the establishment of the Commission on Critical Infrastructure Protection in 1996 by Executive Order 13010, each and every one of these devices and systems has been recognized and acknowledged as critical to the complex interconnectivity to our nation—“Their incapacity or destruction would have a debilitating impact on the defense or economic security of the United States [30].” Consider the potential damage caused by a slight alteration of a single line of code for an AI system controlling the gates of the Hoover Dam. The effects of 10 trillion gallons of water suddenly released from Lake Meade would result in a horrific catastrophe [31].

Adversaries of the U.S. designate all critical infrastructure as high value targets, drawing their constant attention and interest from a technical targeting perspective. Defending the security systems, public-facing networks, and internal industrial controls is paramount to the life and safety of the people in the vicinity of such critical infrastructure. Moreover, this defense is vital to our nation as a whole, through our national strategy. The use of AI for biometrics, HDS, and CIP brings both promising innovation and efficiency as well as potential new risks. The key will be to maximize the benefits of AI while minimizing risks and mitigating complicated safety and security challenges.

References

1. Department of Homeland Security. (n.d.). Rapid screening tool: The AVATAR . Retrieved from <https://www.dhs.gov/sites/default/files/publications/Rapid%20Screening%20Tool-NCBSI-AVATAR-Jan2014.pdf>
2. Department of Homeland Security. (2017, July 11). *Critical infrastructure sectors*. Retrieved from <https://www.dhs.gov/critical-infrastructure-sectors>
3. Reese, S. (2013, January 8). *Defining Homeland Security: Analysis and Congressional considerations*. Congressional Research Service. Retrieved from <https://fas.org/sgp/crs/homesecc/R42462.pdf>
4. Ferguson, W., Diller, D., Moffitt, K., Roberts, B., Coon, W., Mateas, M., ... & Tannenbaum, K. (2015). *Immerse: Interactive mentoring for multimodal experiences in realistic social encounters*. Cambridge, MA: Raytheon Corporation. Retrieved from <http://www.dtic.mil/docs/citations/ADA625663>
5. Sallach, D. L., & North, M. J. (2010). Multigame dynamics: Structures and strategies. In I. Bosse, A. Geller, & C. M. Jonker, (Eds.), *Multi-agent-based simulation XI*. (pp. 108–120). New York: Springer.
6. Ross, K. G., Klein, G. A., Thunholm, P., Schmitt, J. F., & Baxter, H. C. (2004, July–August). The recognition-primed decision model. *Military Review*, 6–10. Retrieved from <http://www.au.af.mil/au/awc/awcgate/milreview/ross.pdf>
7. Hodgins, C. R. (2014). *AlphaACT*. Westminster, CO: AlphaTRAC. Retrieved from <https://www.alphatrac.com/alphaact.aspx>
8. McNamara, L. A., Trucano, T. G., & Gieseler, C. (2011). *Challenges in computational social modelling and simulation for national security decision making*. Albuquerque, NM: Sandia National Laboratories. Retrieved from <http://www.dtic.mil/docs/citations/ADA555381>
9. Doare, R., Danet, D., Hanon, J., & de Boisboissel, G. , Eds. (2014). *Robots on the battlefield: Contemporary perspectives and implications for the future*. Combat Studies Institute Press.
10. Roff, H.M., & Moyes, R. (n.d.). Project: Artificial Intelligence, Autonomous Weapons, and Meaningful Human Control. Discussed within *Meaningful Human Control or Appropriate Human Judgment? The Necessary Limits on Autonomous Weapons* briefing paper presented at the Informal Meeting of Experts on Lethal Autonomous Weapons Systems, UN Convention on Certain Conventional Weapons, The United Nations Office at Geneva, April 11, 2016. Retrieved from <https://globalsecurity.asu.edu/robotics-autonomy>
11. Work, B. (2015, December 14). CNAS Defense Forum. *U.S. Department of Defense*. Retrieved from <https://www.defense.gov/News/Speeches/Speech-View/Article/634214/cnas-defense-forum/>
12. Pellerin, C. (2016, October 31). Deputy Secretary: Third Offset Strategy bolsters America’s military deterrence. *U.S. Department of Defense*. Retrieved from <http://www.defense.gov/News/Article/Article/991434/deputy-secretary-third-offset-strategy-bolsters-americas-military-deterrence?source=GovDelivery>

13. Lei, Z. (2016, August 19). Nation's next generation of missiles to be highly flexible. In *China Daily*. Retrieved from http://www.chinadaily.com.cn/china/2016-08/19/content_26530461.htm
14. Litovkin, D., & Litovkin, N. (2017, May 31). Russia's digital doomsday weapons: Robots prepare for war. *Russia beyond the headlines*. Retrieved from https://www.rbth.com/defence/2017/05/31/russias-digital-weapons-robots-and-artificial-intelligence-prepare-for-wa_773677
15. Mizokami, K. (2016, February 25). The Navy's new AI missile sinks ships the smart way. *Popular Mechanics*. Retrieved from <http://www.popularmechanics.com/military/weapons/a19624/the-navys-new-missile-sinks-ships-the-smart-way/>
16. O'Connor, T. (2017, April 19). Russia built a robot that can shoot guns and travel to space. *Newsweek*. Retrieved from <http://www.newsweek.com/russia-built-robot-can-shoot-guns-and-travel-space-586544>
17. Keller, J. (2010, March 1). Air Force to use artificial intelligence and other advanced data processing to hit enemy where it hurts. *Military and Air Force Electronics*, 21(3). Retrieved from <http://www.militaryaerospace.com/articles/print/volume-21/issue-3/news/news/air-force-to-use-artificial-intelligence-and-other-advanced-data-processing-to-hit-the-enemy-where-it-hurts.html>
18. Freedberg, S. J. (2017, July 13). Artificial intelligence will help hunt Daesh by December. *Breaking Defense*. Retrieved from <http://breakingdefense.com/2017/07/artificial-intelligence-will-help-hunt-daesh-by-december/>
19. Ulery, B. T., Hicklin, R. A., Buscaglia, J., & Roberts, M. A. (2011). Accuracy and reliability of forensic latent fingerprint decisions. *Proceedings of the National Academy of Sciences*, 108, 7733–7738.
20. Zhou, Y., Pandey, R., & Govindaraju, V. (2015, June). Deep secure encoding: An application to face recognition. *arXiv:1506.04340*
21. Hansen, J. H. (2015, October). *Robust speech processing & recognition: Speaker ID, language ID, speech recognition/keyword spotting, diarization/co-channel/environmental characterization, speaker state assessment*. Richardson, TX: Texas University at Dallas. Retrieved from <http://www.dtic.mil/docs/citations/ADA623029>
22. Frischholz, R. W., & Dieckmann, U. (2000). BioID: A multimodal biometric identification system. *Computer*, 33, 64–68.
23. Cain, L. R. (2014). *Feature selection on hyperspectral data for dismount skin analysis*. (Unpublished thesis). Air Force Institute of Technology, Wright-Patterson Air Force Base, OH.
24. Fleming, S. B. (2014, March). *Identification of a smartphone user via keystroke analysis*. Monterey, CA: Naval Postgraduate School.
25. Chihaoui, M., Elkefi, A., Bellil, W., & Amar, C. B. (2016). A survey of 2D face recognition techniques. *Computers*, 5(4), 21. doi:10.3390/computers5040021
26. Culhane, M. (2014, April 1). *Survey of technologies for the airport border of the future*. Ottawa, Ontario, Canada. National Research Council of Canada. Retrieved from <http://www.dtic.mil/docs/citations/AD1016982>

27. Skouson, M. B. (2007). *How the Air Force should stay engaged in computer vision technology development*. Maxwell AFB, AL: Center for Strategy and Technology, Air War College.
28. Parkin, S. (2015, July 16). Killer robots: The soldiers that never sleep. *BBC*. Retrieved from <http://www.bbc.com/future/story/20150715-killer-robots-the-soldiers-that-never-sleep>
29. DHS. (2017, July). *Narrative analysis: Artificial intelligence*. Retrieved from https://www.oodaloop.com/wp-content/uploads/2017/07/OCIA-NetAssessment-Artificial_Intelligence_Narrative_Analysis.pdf
30. Exec. Order No. 13010, 3 C.F.R. 37347 (1996).
31. Scott, S. (2015, September 21). How much damage would be done by the Hoover Dam breaking open?. *Quora*. Retrieved from <https://www.quora.com/How-much-damage-would-be-done-by-the-Hoover-Dam-breaking-open>

4

Chemical, Biological, Radiological, and Nuclear Defense/Weapons of Mass Destruction

Joel Hewett and Gregory Nichols, HDIAC

U.S. warfighters have already deployed AI and machine learning technologies on the battlefield—or, rather, approximately 10,000 feet above it. In December 2017, DoD’s newly-commissioned Algorithmic Warfare Cross-Functional Team successfully deployed an AI-enabled tool in an operational theater for the first time in American military history. The effort, known as Project Maven, applied an object recognition algorithm to the live video feed of unmanned aircraft systems engaged in surveillance during Operation Inherent Resolve to analyze the drone’s reams of footage to generate actionable intelligence [1].

AI promises to drastically reduce the amount of time and effort needed to perform “tedious work,” such as “low-level counting activity” and data entry, which will allow for quicker analysis of still-image and video stream datasets to identify major threats [1]. Moreover, prominent defense analysts have highlighted imagery analysis as one of AI’s most “remarkable capabilities” suited for near-term military use [2].

AI-enabled imagery analysis is likely to be used for the detection and monitoring of CBRN and WMD threats and assets. Former Deputy Defense Secretary Robert Work, who established Project Maven, said in November 2015 these programs could soon be used to efficiently monitor worldwide WMD sites, triggering an alert when a “facility is moving from a benign posture to a threatening posture [3].” Such monitoring may also be further tailored to search for specific assets types, such as “a [ballistic missile] transporter that’s 15 meters long, 4.7 meters wide,” Work offered as an example [3].

While at the time of writing Project Maven is not explicitly focused on CBRN or WMD threats, its status as DoD’s first field deployment of an AI-enabled warfighting tool—and the unique and critical role that imagery analysis plays in the CBRN Defense and Countering WMD missions [4,5]—make it an exemplar of the state-of-the-art in this subfield. The parameters of Project Maven closely reflect how AI and machine learning algorithms are most likely to be applied by DoD to the CBRN/WMD defense mission over the next 5 to 10 years. This general point is supported by two conclusions drawn from HDIAC’s technical analysis of the project, both of which are elaborated on in the next section:

1. AI-enabled imagery analysis software for CBRN/WMD applications is years—if not a decade or more—away from functioning in the “plug-and-play” mode desired by civilian as well as uniformed DoD leaders.
2. Future deployments of AI-enabled algorithms for imagery analysis are all but certain to grow increasingly processor-intensive, suggesting that machine learning engineers must codesign AI software and supporting hardware *simultaneously* in order to deliver workable applications useful to DoD activities.

Project Maven

Ultimately, technologically-minded strategic thinkers and military leaders hope to apply the immense computing power of AI-based systems to sensing systems in a modular, self-contained configuration. Indeed, when speaking to *Seapower* magazine in May 2016, CPT Jeffrey J. Czerewko, now Commander, Naval Air Force, U.S. Pacific Fleet, described his vision for Naval AI with exactly that tenor of language. He described ideal field-deployable machine learning systems as “platform and domain agnostic” in their use, capable of being hooked up to any existing sensor system as “little black boxes” whose internal workings will require no specific maintenance or continual monitoring [6].

That vision remains powerful, and—it should be noted—fully viable. However, through the end of 2017, Project Maven has demonstrated the infeasibility of delivering this advanced type of “plug-and-play” capability for American warfighting use in the medium-term.

First, the construction and coding of Project Maven’s foundational object identification algorithm was labor-intensive and slow. Program staff manually identified and labeled more than 150,000 images simply to build the starting dataset from which the AI algorithm would teach itself to analyze drone footage [1]. Continued use of the imagery analysis system beyond its initial deployment date will require the manual cataloging of as many as 750,000 *additional* images into the dataset. This will be performed to:

- optimize the algorithm’s performance to match the distinct attributes of its specific desert environment [1]
- decrease the likelihood of the system committing misidentifications due to noisy or substandard image quality [7]
- defend it from disruptive or adversarial attempts to trick or spoof its computer vision [8]

The same requirements for operational maintenance will apply to AI-enabled systems implemented in the near-term, until more general-purpose technological imagery pre-processing programs can be developed to shorten the period needed to train an AI-enabled machine program to begin teaching itself. Indeed, CNA recently noted that “little attention” has been paid to the “amount and kind of data necessary to train military AI systems” [9].

Second, despite the immense amount of computer processing power available to DoD entities (to tap or directly procure), Project Maven’s AI-enabled analytical code once compiled required the assembly of custom-built IT hardware to support its operation [1]. Even when the AI was training itself on only the base dataset of 150,000 images, its operations-per-second load was so processor- and graphics-intensive that it was deemed unable to run off pre-existing DoD networks, including those that were fully cloud-based [1]. The requirement of bespoke coding and engineering labor in operating an AI-enabled tool *for a single video feed* is not a workable proposition beyond the immediate near-term.

This challenge is a novel one even for Project Maven’s state-of-the-art system architecture. Previous instances of AI-enabled imagery analysis systems, typically

ensconced in the laboratories of R1 academic research institutions and highly-capitalized private enterprises, have simply drawn on brute force processing resources to operate. When Google designed an AI-equipped neural network in 2012 to sort through millions of YouTube videos and images to recognize cat pictures, the search company's team simply assembled an array of 16,000 standard-run, commercially-available central processing units (CPUs) to perform the feat [10].

However, future iterations of the Google AI system (renamed Google Brain shortly thereafter) butted heads against fundamental limitations posed by current computing hardware. During 2016 and 2017, Google engineers turned to an emerging, little-known type of processing chip to boost their network's nameplate processing capacity. Tensor processing units (TPU) promise to enable AI-system-wide performance improvements as much as 30 times greater than previous chips [11]. In mid-2017, Google installed an even more powerful version of the initial TPU model with positive results [12].

Notably, even given the relative simplicity of Google's unplug-and-replace operation (removing CPU and graphics processing units to make room for the tensor units), the effort was not a simple upgrade job. The TPUs required many months of debugging and software-hardware integration work before they could even function [13].

CBRN/WMD Imagery Analysis

As the 2017 *National Security Strategy of the United States of America* stresses, the detection of WMD assets remains a key priority action for the nation's defense community. The strategy states that the United States will "better integrate intelligence...operations to ensure that frontline defenders have the right information and capabilities to respond to WMD threats" [14]. Elsewhere, the Army has called upon the scientific and technical enterprise of the nation to help provide next-generation "environmental monitoring and threat warning services" for DoD's global ballistic missile defense mission [15]. Therefore, delivering capabilities to achieve this mission has been of high interest to AI researchers and engineers.

In July 2017, researchers at the University of Missouri reported the results of a landmark effort to apply deep convolutional neural networks to the analysis of Chinese surface-to-air missile (SAM) launch sites. (These sites display a distinctive configuration, typically characterized by the presence of four to six small launch pads, set in a broad symmetrical and circular shape [16].)

Drawing upon a training dataset of around 1.6 million open-source images of known SAM sites, the researchers restricted the AI program's access to no more than 100 such positive or affirmative image examples from which it could learn. They then directed the program to conduct a broad-area survey of satellite imagery collected over 55,078 square miles of terrain in southeastern China. The study trained multiple AI-enabled algorithms to auto-select potential candidate sites and prioritize them for imagery personnel to analyze (see Figure 1).




Figure 1. Web user interface used for human review of AI-identified and ranked candidate SAM sites. The image on the left presents a contextual view of a candidate site, while the image on the right offers a zoomed-in view of the site, providing the viewer with additional detail for inspection. This candidate was correctly identified as a SAM site [16]. (Released)

By doing so, the system reduced the search area requiring human review to only 0.15 percent of the 55,078-square-mile search area. Moreover, this process, as the authors explain, allowed “four novice imagery analysts with no prior imagery analysis experience” to complete the site search nearly 81 times “faster than a traditional visual search [conducted] over an equivalent land area...while achieving nearly identical statistical accuracy” [16].

As the University of Missouri research effort, Project Maven, and Google Brain make clear, an AI-enabled image analysis algorithm can train itself to detect and identify nearly any type of target, provided that a sufficiently large training dataset exists. In other words: replace the exceptionally small training set of SAM sites (<100 images) with images of chemical weapons production facilities, biological weapons testing sites, or radioactive material extraction fields—and AI-enabled programs will significantly advance the CBRN/WMD defense missions of the United States.

Emerging Hardware Considerations

As the experience of Project Maven indicates, hardware and processing power access considerations are likely to place limitations on black box-like applications of AI-enabled algorithms for imagery analysis in the medium-term. This issue is apparent in the SAM identification study performed at the University of Missouri: when using the reduced final training set of <100 images, the *total* processing time for the neural network program was roughly 23 hours, even when performed on a high-powered commercial-grade desktop computer station with a high-rate connection to network attached storage [16].

As Google's experience with TPUs makes clear, the path forward to achieving AI-enabled analysis faster and more generalized lies, in part, with the development of hardware specifically designed and optimized for AI use. Multiple leading firms, including Nvidia, Intel, Qualcomm, ARM, and Microsoft are indeed seeking to do just that [17,18]. In the far-term, the usefulness of AI in the analysis of massive datasets may depend on successful yet-to-be-determined developments in quantum computing and/or neuromorphic spiking neural nets [19]. As the Defense Science Board noted in 2016, "technology advances in high-throughput, embedded processing, and machine learning offer the promise of onboard processing of high-resolution, multi-source airborne ISR sensor data" [20].

In order to place AI-enabled imagery analysis at the edge of deployed platforms—or as close to the field as possible—future systems will need to be designed with both hardware and software components in mind. DoD has already taken some steps in this direction [21].

R&D removed from the operational side of AI-enabled systems has recently underscored this future path of development. In December 2017, after conducting a comprehensive review of AI-enabled uses of deep-learning neural networks, a team of engineers and entrepreneurs from MIT and Nvidia cautioned that the continued design of AI systems untethered to specific hardware considerations would reduce their accuracy and efficiency. Instead, systems engineers should codesign algorithm models and hardware contemporaneously, to both "maximize accuracy and throughput, while minimizing energy and cost" [22].

Conclusion

DoD leaders have stated that the most promising yet untapped datasets collected by their forces are image- and/or video-based. As technology and national security analyst Gregory Allen explains:

Every day, US spy planes and satellites collect more raw data than the Defense Department could analyze even if its whole workforce spent their entire lives on it. [...] Unfortunately, most of the imagery analysis involves tedious work—people look at screens to count cars, individuals, or activities, and then type their counts into a PowerPoint presentation or Excel spreadsheet [1].

"Worse," Allen concludes, "most of the sensor data just disappears—it's never looked at—even though the department has been hiring analysts as fast as it can for years [1]."

Plug-and-play applications of AI in the CBRN/WMD detection space remain years away from operational use. For DoD to tap into the wealth of imagery sensor data it collects every day, future efforts in this subfield must: (a) investigate ways of streamlining algorithm coding and self-training procedures and (b) work in tandem with hardware manufacturers to codesign machine learning hardware-software systems that optimize the function of both.

References

1. Allen, G. C. (2017, December 21). Project Maven brings AI to the fight against ISIS. *Bulletin of the Atomic Scientists*. Retrieved from <https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374>
2. Allen, G., Chan, T. (2017, July). *Artificial Intelligence and National Security* (Rep.). Cambridge, MA: Harvard University, Harvard Kennedy School, the Belfer Center for Science and International Affairs. Retrieved from <https://www.belfercenter.org/sites/default/files/files/publication/AI%20NatSec%20-%20final.pdf>
3. Weisgerber, M. (2016, April 28). The increasingly automated hunt for mobile missile launchers. DefenseOne. Retrieved from <https://www.defenseone.com/technology/2016/04/increasingly-automated-hunt-mobile-missile-launchers/127864>
4. Gwertzman, B. (1979, October 26). U.S. monitors signs of atom explosion near South Africa. *New York Times*. Retrieved from <https://www.nytimes.com/1979/10/26/archives/us-monitors-signs-of-atom-explosion-near-south-africa-seeks-more.html>
5. Park, J., & Helmboldt, J. (2017). Detection of large scale man-made explosions by the Global Navigation Satellite System. *Journal of the Homeland Defense & Security Information Analysis Center*, 4(2), 24-29. Retrieved from https://www.hdiac.org/wp-content/uploads/files/journals/Journal_Volume_4_Issue_2_Final/HDIAC%20Journal_Volume%204_Issue%202_FINAL%20WEB.html
6. Matthews, W. (2016, May 18). Navy's Minotaur system is a step toward automated data analysis. *Seapower*. Navy League of the United States. Retrieved from <http://seapowermagazine.org/stories/20160518-data.html>
7. Hosseini, H., Ziao, B., & Poovendran, R. (2017, July 20). *Google's cloud vision API is not robust to noise*. Paper presented at the 16th IEEE International Conference on Machine Learning and Applications, Cancun, Mexico. Retrieved from <https://arxiv.org/pdf/1704.05051.pdf>
8. Kurakin, A., Goodfellow, I. J., & Bengio, S. (2016, July 8). *Adversarial examples in the physical world* (Tech. Rep.). Google, Inc. Retrieved from <https://arxiv.org/pdf/1607.02533v1.pdf>
9. Iliachinski, A. (2017, October.) *Artificial intelligence & autonomy: Opportunities and challenges* (Rep.). Arlington, VA: CNA. Retrieved from https://www.cna.org/CNA_files/PDF/DIS-2017-U-016388-Final.pdf
10. Markoff, J. (2012, June 25). How many computers to identify a cat? 16,000. *New York Times*. Retrieved from <https://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html>
11. Sato, K., Young, C., & Patterson, D. (2017, May 12). An in-depth look at Google's first tensor processing unit (TPU) [Web log post]. Google Cloud Big Data and Machine Learning Blog. Retrieved from <https://cloud.google.com/blog/big-data/2017/05/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu>
12. Simonite, T. (2017, July 25). The rise of AI is forcing Google and Microsoft to become chipmakers. *Wired*. Retrieved from <https://www.wired.com/story/the-rise-of-ai-is-forcing-google-and-microsoft-to-become-chipmakers/>

13. Lewis-Kraus, G. (2016, December 14). The great A.I. awakening: How Google used artificial intelligence to transform Google Translate, one of its more popular services—and how machine learning is poised to reinvent computing itself. *New York Times Magazine*. Retrieved from <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>
14. The White House. (2017, December). *National Security Strategy of the United States of America*. Retrieved from <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf>
15. U.S. Army. (2017, September 22). *Warfighters' Science and Technology Needs Bulletin*. Fort Eustis, VA: U.S. Army Training and Doctrine Command. Retrieved from http://www.arcic.army.mil/APP_Documents/Army-Warfighters-ST-Needs-Bulletin.pdf
16. Marcum, R. A., Davis, C. H., Scott, G. J., & Nivin, T. W. (2017, October 6). Rapid broad area search and detection of Chinese surface-to-air missile sites using deep convolutional neural networks. *Journal of Applied Remote Sensing*, 11(4). doi:10.1117/1.JRS.11.042614
17. Melendez, S. (2016, December 7). Chipmakers are racing to build hardware for artificial intelligence. Fast Company. Retrieved from <https://www.fastcompany.com/3065843/artificial-intelligence-chips-race-vidia-intel>
18. Versace, M. (2017, October 20). Does artificial intelligence require specialized processors? The New Stack. Retrieved from <https://thenewstack.io/ai-hardware-software-dilemma/>
19. Vorhies, W. (2017, November 14). The three-way race to the future of AI. Quantum vs. neuromorphic vs. high performance computing [Web log post]. Data Science Central. Retrieved from <https://www.datasciencecentral.com/profiles/blogs/the-three-way-race-to-the-future-of-ai-quantum-vs-neuromorphic-vs>
20. Defense Science Board. (2016, July). *Report of the Defense Science Board Summer Study on Autonomy*. Department of Defense: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics. Retrieved from <https://www.hsdl.org/?view&did=794641>
21. Weisberger, M. (2017, December 21). The Pentagon's new artificial intelligence is already hunting terrorists. DefenseOne. Retrieved from <https://www.defenseone.com/technology/2017/12/pentagons-new-artificial-intelligence-already-hunting-terrorists/144742/>
22. Sze, V., Chen, Y-H., Yang, T-J., & Emer, J. S. (2017, December). Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12), 2295-2329. doi:10.1109/JPROC.2017.2761740

5

Artificial Intelligence Applications to Military Medicine

Gopi Deshpande, Auburn University
 Pradyumna Lanka, University of California, Merced
 Susan A. Youngblood, Auburn University

Introduction

The applications of AI in medicine are wide-ranging. Medicine in general, and disease diagnosis in particular, has evolved through the development of precise objective biomarkers for disease detection. For example, diabetes can be definitively diagnosed by the assessment of blood glucose levels. These objective and replicable biomarkers are sparse in certain fields of medicine, such as mental health, which are relevant to the defense community. In fact, one of the primary applications of machine learning, a subfield within AI, is identifying clinically relevant biomarkers that are reproducible and can be used to identify disorders on an individual subject level. Of specific interest to the defense community are diseases such as post-traumatic stress disorder (PTSD) and post-concussion syndrome (PCS), a sequel of mild traumatic brain injury (mTBI), which have high prevalence rates in military veterans [1,2].

The standards required for single-subject predictions that are necessary for automated disease diagnosis are more rigorous than group differences identified from population-level statistics. Therefore, the findings from previous studies identifying group differences in brain function are not directly applicable while searching for biomarkers. This necessitates the use of machine learning in neuroimaging applications [3]. Fortunately, the number of studies exploring single-subject prediction in mental health disorders has increased exponentially in recent years with more than 500 papers published in the past quarter century (see Figure 1) [3]. Apart from diagnostic classification, AI in medicine can also be used for monitoring disease progression and treatment response.




Figure 1. A bar plot showing the number of studies performing single-subject disease predictions over the past 22 years for several neurological and mental health disorders.

Below, we detail the applications of AI, especially machine learning in mental health disorders and the available opportunities in emerging fields, such as multimodal data fusion, deep learning in neural imaging, and data sharing initiatives. In each section, we examine some of the challenges that must be addressed to realize the full potential of these applications.

Applications of AI in Mental Health Disorders and Brain Injuries

Mental health disorders arise out of complex alterations of the brain. Distinct neural signatures have been observed for healthy and clinical populations [4]. AI can aid in diagnosis and prognosis of neurological disorders, including in PTSD, a disorder that develops in individuals exposed to traumatic situations. It can also aid in identifying potential suicide attempters.

Although machine learning has limitations, it can help identify biomarkers for diagnosing spectrum disorders. These disorders show distinct pathologies with multiple etiologies under the umbrella of the same diagnostic category. Spectrum disorders include not only autism spectrum disorder (ASD), schizophrenia, and attention-deficit/hyperactivity disorder (ADHD) but also PTSD. A binary distinction between controls and disease populations ignores the underlying heterogeneity of these disorders in the general population.

Machine learning, especially unsupervised learning (in which the machine learning algorithms clusters data, unguided, based on similarities they identify), can help with the sub-categorizations of such disorders by identifying novel sub-types of the disease and improving sub-groupings in a data-driven way. For example, data-driven groupings of subjects are reported to be superior to conventional diagnostic groupings [5,6]. Studies using supervised machine learning are more limited. In supervised machine learning, in which the algorithm is trained by providing output variables (e.g., ranges within a set of parameters yield a diagnosis of PTSD), the performance of the algorithms is constrained by the clinical labels associated with the data and the ability of current clinical diagnostic instruments to accurately capture the underlying neurobiology of the disorders. Hence, the best possible performance will mimic the disease labels obtained from clinical assessment and is thus constrained by their errors and limitations [7].

A case in point is the diagnostic label associated with spectrum disorders such as ASD, which are highly heterogeneous and have large symptom variability. In such cases, the label might not be entirely informative, thus limiting the identification of reliable biomarkers [3]. Some other issues associated with automated disease diagnosis for clinical applications are high comorbidity of neurological disorders and differential diagnoses of disorders with different etiologies but overlapping symptoms [8]. Examples of disorders with overlapping symptoms include schizophrenia, bipolar, unipolar and mood disorders [8] as well as PTSD and PTSD associated with PCS/mTBI [5].

Unfortunately, despite the serious nature of PTSD, the current methods for diagnosis rely on self-reported symptom scores and clinician observation. PTSD has a high prevalence in combat veterans and high comorbidity with PCS/mTBI. Symptoms of PTSD and PCS overlap, complicating diagnosis and treatment. Because the underlying etiologies appear to be different, though, machine learning

and neuroimaging can help [5]. Similar work has been applied to ADHD diagnoses [9]. Resulting objective assessments can help in making “return to duty” decisions and also prevent malingering and false insurance claims. Such efforts are already underway with Rangaprakash et al. identifying elevated static brain connectivity and reduced dynamic brain connectivity in veterans with PTSD compared to controls [5], with such measures providing predictive power at the individual subject level. It must be noted that such studies must be replicated in larger samples to be useful as a potential biomarker. Progress is also being made in using algorithms to detect repeated mTBI, which is laborious to diagnose manually from MRIs because of the small lesion sizes [10,11].

Suicide occurs at a higher rate in the U.S. military than in the general population, in part because of military age demographics, gender demographics [12], and access to firearms [13]. AI applications to suicide prevention are being employed to refine the results of psychometric questionnaires, better identifying potential suicide attempters. Support vector machines, a variety of supervised machine learning, have refined the results of the Barratt Impulsiveness Scale Version 11 and International Personality Disorder Examination - Screening Questionnaire psychometric tools to better weight questionnaire responses [14].

AI and data mining have also been used to create decision trees psychometric results [15]. In each case, researchers compared suicidal participants to non-suicidal participants to model the weights of risk factors. AI work in suicide prevention relies on psychometric questionnaires, though, and early detection of suicidal ideation may be hindered by active duty personnel's reticence to report ideation [16,17] and thus present psychometric data.

Classification Using Multimodal and Phenotypic Data

Any single imaging modality (e.g., single-modality MRI) may not capture the subtle differences in the underlying neuropathology for all neurological disorders. Combining image-based features from several imaging modalities (e.g., structural MRI and resting-state functional MRI, or rs-fMRI) usually gives better classification performance than any single imaging modality for disease diagnosis [18,19]. Several resting-state functional connectivity (RSFC) studies ignore the richness of the temporal information and use static connectivity measures for classification. Using dynamic connectivity measures that capture the synchronization of brain regions across time have better sensitivity to neurological disorders and result in better accuracy [20]. Combining phenotype information, such as cognitive and behavioral scores and demographic information, along with clinical measurements, often leads to even better performance than studies that utilize imaging-based features alone [7,21-25]. These results stress the importance of adding any measures sensitive to the presence or progression of the disease in the classification model.

Deep Learning in Neuroimaging

Deep learning uses artificial neural networks to model non-linear relationships and classify data. It has been successful in giving state-of-the-art classification accuracies in computer vision, speech recognition, and language processing [26,27]. In 2016, a program called Alpha Go, powered by Google's DeepMind, which utilizes deep learning, defeated the world's top ranking human Go player. DeepMind has already partnered with London hospitals for detection of eye

diseases, such as macular degeneration and cancerous tissues in the head and neck regions [28,29]. Unlike traditional machine learning methods, where the optimal features need to be hand-coded using expert knowledge, deep learning methods use complex non-linear transformations from the raw data to extract optimal features to be used for classification.

Computer vision (CV) has applications in medical imaging for coronary structures, chronic wounds, neuronal dysfunction [30], ophthalmic structures [31], and cancers [32], as well as for training military medical staff without cadavers [33]. CV has benefited from large sample sizes. The sample sizes used in typical CV studies are orders of magnitude larger than what can be acquired from brain imaging. For example, ImageNet, a standard dataset used to benchmark the performance of CV algorithms, has more than a million data samples categorized into 1,000 classes [3]. Further, in many CV studies, the size of the dataset is often augmented by suitable transformation, such as rotation, shear, and scaling, on the images and then the classifier is trained using the original and transformed images, leading to better classification performances.

Deep learning in neuroimaging, on the other hand, has traditionally been challenged by the problem of small sample sizes, limiting the usefulness of such studies. Small sample sizes can cause overfitting, the problem of a model with many parameters not handling a new dataset well, thereby not classifying properly. Regularization strategies can help compensate for the smaller datasets and are commonly part of current deep learning algorithms [7]. In fact, convolutional neural networks, which use weight sharing, have been reported to give some of the best classification performances in neuroimaging applications [7].

Transfer learning is another method that has been used to counteract the limitations typical of neuroimaging datasets (the effects of small sample sizes, as well as different data-acquisition protocols and site variability, which are discussed below). In transfer learning, features learned using one dataset aid in solving machine learning problems in another set of data [7]. For example, transfer learning has been used in the classification of subjects with Alzheimer's disease using initial parameters learned from natural images [34] and another Alzheimer's dataset [35]. Recent studies have reported excellent classification performances with deep learning methods [36-38]. A summary of classification studies using deep learning methods is available in Vieira et al [7].

With the availability of large datasets, complex algorithms, such as deep learning models with several hidden layers, could model the complex relationships between the features by fine-tuning the parameters and capture subtle differences in the brain function between the clinical and healthy populations. These differences previously would have been left undetected by the use of simpler models applied to neuroimaging data. An example illustrating the power of deep nets with multiple hidden layers is shown in Figure 2.




Figure 2. This figure illustrates the increased classification ability of classes with deep networks projected on a 2D map. In the figure, each dot represents a subject. Classes that are not clearly separable with shallow networks (depth 1 and 2) become increasingly separable with a deeper architecture (depth 3) [37]. (Released)

Data Sharing and Big Data Neuroimaging Initiatives

A number of limitations are being addressed by data sharing. As discussed above, perhaps the biggest hurdle hampering the search for clinically relevant neuroimaging biomarkers is small sample sizes, such as samples collected from a single acquisition site. High accuracies reported in studies with small sample sizes, which do not sample adequately from the entire population distribution, might not generalize to larger heterogeneous disease populations. Further challenges preventing identification of reliable biomarkers include small effect sizes, low disease prevalence, and clinical and data heterogeneity. Even age and disease maturation can vary imaging biomarkers because of the plasticity and dynamic adaptability of the brain [4,39].

Data sharing and big-data initiatives offer significant advantages for the application of machine learning algorithms in neuroimaging studies. Large sample sizes capture the disease heterogeneity in the clinical populations, which is especially useful for generalization of the results to new subjects. Results of studies with large sample sizes and effect sizes are a first step in the translation of key findings in research studies to the clinical setting. Big data initiatives promote exploratory data analyses, compared to hypothesis-driven data analysis traditionally used in neuroimaging studies [8]. Large-scale data sharing initiatives also promote the principles of open science, and results obtained on such large datasets help in reproducibility.

Data sharing initiatives hold the promise of providing an objective estimation of performance (benchmarking), assessing current classification, and improving classification, but must first overcome the lack of standardization. A key problem is inter-site variability, such as variability in data acquisition protocols and demographic and experimental factors, which Colby et al. [40] note limits the usefulness of such multisite studies. Currently, single-site studies in which the data from the same acquisition site is used for training and testing report much better classification performances than when they are from distinct image acquisition sites [40,41]. Studies using large heterogeneous samples report lower performances [3].

Pooling the data from poorly coordinated retrospective multisite studies can yield poor neuroimaging data. For example, in the ADHD-200 global competition, researchers from the University of Alberta achieved a higher accuracy (62.5

percent) using just phenotypic data than using neuroimaging-based metrics [42]. Continued efforts to share data should result in standardizing neuroimaging protocols and coordinating acquisition protocols—which Milham et al. refer to as “harmonizing” acquisition protocols—across imaging sites [8], making the resulting classifications more clinically useful. In addition, data sharing allows researchers to compare classification performances obtained by data acquired from different data processing pipelines and classification models, which is currently impossible, given that traditional studies use different datasets.

Several data-sharing efforts for both healthy and clinical populations are complete, currently underway, or planned and are extending the range of specialists tackling disease classification problems, especially pertaining to mental health. Examples of such initiatives include the 1000 Functional Connectomes Project [43]; International Neuroimaging Data-sharing Initiative [44]; OpenfMRI [45]; U.K. Biobank [46]; Human Connectome Project [47]; Autism Brain Imaging Data Exchange [48]; ADHD-200 [22]; Alzheimer's Disease Neuroimaging Initiative [49]; Brain Genomics Superstruct [50]; and the Pediatric Imaging, Neurocognition, and Genetics repository [51]. Machine learning contests such as the ADHD-200 Global Competition and the Alzheimer's Disease Big Data DREAM Challenge and Brain Hack help people from diverse research backgrounds work on challenging problems in neuroimaging, and the availability of preprocessed datasets (e.g., The Preprocessed Connectomes Project [52]) have reduced their barrier to entry even further [8].

With a large amount of data collected from multiple sites, there is renewed optimism in the search for reliable neuroimaging-based disease biomarkers that generalize well to the larger population. Big data analyses in neuroimaging are expected to accelerate with advancements in technology (particularly in storage and computational power in both traditional and cloud computing), software that can run in parallel on graphical processing units, and advances in our understanding of machine learning and deep learning algorithms.

Additional Challenges

Cross-validation accuracy in multisite studies can also be an unreliable estimator of the true predictive ability of the classifiers [53,54] due to hidden correlations in training and the validation data. In addition, the applicability of automated diagnostic tools limited for the general population could cause false positives due to low disease prevalence in the general population. Furthermore, head motion in neuroimaging adds spurious noise, which can significantly affect the validity of the results obtained from the studies. Several studies have shown that head motion alters RSFC based metrics in a distance-dependent way [55,56] can significantly alter the classification accuracy in distinguishing controls from hyperkinetic disease populations. Fair et al. [57], though, report that thorough motion correction strategies could improve the classification accuracy.

Conclusion

AI in military-relevant medical applications has progressed significantly, particularly in the domain of making diagnoses once reliant on patient responses and clinician observation (e.g., PTSD and suicide risk) more objective and in using clinically relevant biomarkers to monitor both disease progression and patient

responses to treatment. AI is allowing the medical field to tease out the underlying causes of conditions with similar symptoms but different etiologies. The strides the field is making in neuroimaging are promising, and these are due in no small part to moving beyond single imaging modalities to take advantage of richer sets of data, to employing deep learning, and to overcoming small sample sizes through data sharing initiatives.

Moreover, the use of AI in medicine can help usher in an era of precision and personalized medicine. Precision medicine aims to tailor the treatment to an individual or a group of individuals by stratifying them based on genetic, immunological, and imaging profiles to provide the best possible treatment response. By identifying the disease diagnosis, subtypes, disease progression and treatment response, the goals of precision medicine can be achieved [8].

Additional progress must be made in standardizing both neuroimaging protocols and data acquisition protocols to achieve the field's goals of reliable single-subject predictions, but researchers are actively pursuing those advances. These tasks are challenging, but current data-sharing initiatives are giving specialists from diverse fields an opportunity to join in this work. With interdisciplinarity comes new perspectives, and blended expertise often ushers in new solutions.

References

1. Tanielian, T., & Jaycox, L. (2008). *Invisible Wounds of War: Psychological and Cognitive Injuries, Their Consequences, and Services to Assist Recovery*. Santa Monica, CA: RAND Corporation. Retrieved from http://www.rand.org/content/dam/rand/pubs/monographs/2008/RAND_MG720.pdf
2. Kang, H., Natelson, B., Mahan, C., Lee, K., & Murphy, F. (2003). Post-traumatic stress disorder and chronic fatigue syndrome-like illness among Gulf War veterans: A population-based survey of 30,000 veterans. *American Journal of Epidemiology*, *157*(2), 141–148. doi:10.1093/aje/kwf187
3. Arbabshirani, M., Plis, S., Sui, J., & Calhoun, V. (2016). Single subject prediction of brain disorders in neuroimaging: Promises and pitfalls. *NeuroImage*, *145*, 137–165. <https://doi.org/10.1016/j.neuroimage.2016.02.079>
4. Horwitz, B., & Rowe, J. (2011). Functional biomarkers for neurodegenerative disorders based on the network paradigm. *Progress in Neurobiology*, *95*(4), 505–509. doi:10.1016/j.pneurobio.2011.07.005
5. Rangaprakash, D., Deshpande, G., Daniel, T. A., Goodman, A. M., Robinson, J. L., Salibi, N., . . . Dretsch, M. N. (2017). Compromised hippocampus-striatum pathway as a potential imaging biomarker of mild-traumatic brain injury and posttraumatic stress disorder. *Human Brain Mapping*, *38*(6), 2843–2864. doi:10.1002/hbm.23551
6. Brodersen, K. H., Deserno, L., Schlagenhaut, F., Lin, Z., Penny, W. D., Buhmann, J. M., & Stephan, K. E. (2014). Dissecting psychiatric spectrum disorders by generative embedding. *NeuroImage: Clinical*, *4*, 98–111. doi:10.1016/j.nicl.2013.11.002
7. Vieira, S., Pinaya, W. H., & Mechelli, A. (2017). Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: Methods and applications. *Neuroscience & Biobehavioral Reviews*, *74*, 58–75. doi:10.1016/j.neubiorev.2017.01.002
8. Milham, M. P., Craddock, R. C., & Klein, A. (2017). Clinically useful brain imaging for neuropsychiatry: How can we get there? *Depression and Anxiety*, *34*(7), 578–587. doi:10.1002/da.22627
9. Bruchmüller, K., Margraf, J., & Schneider, S. (2012). Is ADHD diagnosed in accord with diagnostic criteria? Overdiagnosis and influence of client gender on diagnosis. *Journal of Consulting and Clinical Psychology*, *80*(1), 128–138. doi:10.1037/a0026582
10. Bianchi, A., Bhanu, B., Donovan, V., & Obenaus, A. (2014). Visual and contextual modeling for the detection of repeated mild traumatic brain injury. *IEEE Transactions on Medical Imaging*, *33*(1), 11–22. doi:10.1109/TMI.2013.2269317
11. Bianchi, A., Bhanu, B., & Obenaus, A. (2015). Dynamic low-level context for the detection of mild traumatic brain injury. *IEEE Transactions on Biomedical Engineering*, *62*(1), 145–153. doi:10.1109/TBME.2014.2342653

12. Svan, J. H. (2016, February 9). DOD: Among services, Army had highest suicide rate in 2014. Retrieved from <https://www.stripes.com/news/dod-among-services-army-had-highest-suicide-rate-in-2014-1.393144#.WXjAy9PyuRs>
13. Khazem, L. R., Houtsma, C., Gratz, K. L., Tull, M. T., Green, B. A., & Anestis, M. D. (2016). Firearms matter: The moderating role of firearm storage in the association between current suicidal ideation and likelihood of future suicide attempts among United States military personnel. *Military Psychology, 28*(1), 25–33. doi:10.1037/mil0000099
14. Delgado-Gomez, D., Blasco-Fontecilla, H., Alegria, A. A., Legido-Gil, T., Artes-Rodriguez, A., & Baca-Garcia, E. (2011). Improving the accuracy of suicide attempter classification. *Artificial Intelligence in Medicine, 52*(3), 165–168. doi:10.1016/j.artmed.2011.05.004
15. Morales, S., Barros, J., Echávarri, O., Garcia, F., Osses, A., Moya, C., . . . Tomicic, A. (2017). Acute mental discomfort associated with suicide behavior in a clinical sample of patients with affective disorders: Ascertainning critical variables using artificial intelligence tools. *Frontiers in Psychiatry, 8*. doi:10.3389/fpsy.2017.00007
16. Anestis, M. D., & Green, B. A. (2015). The impact of varying levels of confidentiality on disclosure of suicidal thoughts in a sample of United States National Guard personnel. *Journal of Clinical Psychology, 71*(10), 1023–1030. doi:10.1002/jclp.22198
17. Blocker, G. M., & Miller, J. A. (2013). Unintended consequences: Stigma and suicide prevention efforts. *Military Medicine, 178*(5), 473. doi:10.7205/MILMED-D-13-00056
18. Zhang, Q., Wu, Q., Zhu, H., He, L., Huang, H., Zhang, J., & Zhang, W. (2016). Multimodal MRI-based classification of trauma survivors with and without post-traumatic stress disorder. *Frontiers in Neuroscience, 10*, 292. doi:10.3389/fnins.2016.00292
19. Li, F., Tran, L., Thung, K-H., Ji, S., Shen, D., & Li, J. (2015). Robust deep learning for improved classification of AD/MCI patients. In *Machine Learning in Medical Imaging 5th International Workshop, MLMI 2014, Held in Conjunction with MICCAI 2014, Boston, MA, USA, September 14, 2014. Proceedings* (pp. 240-247). Switzerland: Springer International Publishing.
20. Jin, C., Jia, H., Lanka, P., Rangaprakash, D., Li, L., Liu, T., . . . Deshpande, G. (2017). Dynamic brain connectivity is a better predictor of PTSD than static connectivity. *Human Brain Mapping, 38*(9), 4479-4496. doi:10.1002/hbm.23676
21. Suk, H. I., Lee, S. W., & Shen, D. (2014). Hierarchical feature representation and multimodal fusion with deep learning for AD/MCI diagnosis. *NeuroImage, 101*, 569–582. doi:10.1016/j.neuroimage.2014.06.077
22. The ADHD-200 Consortium. (2012). The ADHD-200 Consortium: A model to advance the translational potential of neuroimaging in clinical neuroscience. *Frontiers in Systems Neuroscience, 6*. doi:10.3389/fnsys.2012.00062
23. Willette, A. A., Calhoun, V. D., Egan, J. M., & Kapogiannis, D. (2014). Prognostic classification of mild cognitive impairment and Alzheimer's disease: MRI independent component analysis. *Psychiatry Research: Neuroimaging, 224*(2), 81–88. doi:10.1016/j.pscychresns.2014.08.005

24. Moradi, E., Pepe, A., Gaser, C., Huttunen, H., & Tohka, J. (2015). Machine learning framework for early MRI-based Alzheimer's conversion prediction in MCI subjects. *NeuroImage*, *104*, 398–412. doi:10.1016/j.neuroimage.2014.10.002
25. Zhang, D., & Shen, D. (2012). Predicting future clinical changes of MCI patients using longitudinal and multimodal biomarkers. *PLoS ONE*, *7*(3), e33182. doi:10.1371/journal.pone.0033182
26. Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). Imagenet classification with deep convolutional neural networks. *Proceedings of the 25th International Conference on Neural Information Processing Systems*, 1097–1105. Lake Tahoe, Nevada.
27. Le, Q. V., Ranzato, M., Monga, R., Devin, M., Chen, K., Corrado, G. S., . . . Ng, A. Y. (2012). Building high-level features using large scale unsupervised learning. *Proceedings of the 29th International Conference on Machine Learning*, 81–88. Edinburgh, Scotland.
28. Baraniuk, C. (2016, August 31). Google DeepMind targets NHS head and neck cancer treatment. BBC News. Retrieved from <http://www.bbc.com/news/technology-37230806>
29. Baraniuk, C. (2016, July 5). Google's DeepMind to peek at NHS eye scans for disease analysis. BBC News. Retrieved from <http://www.bbc.com/news/technology-36713308>
30. Martis, R. J., Lin, H., & Gurupur, V. P. (2017). Frontiers in development of intelligent applications for medical imaging processing and computer vision. *Computers in Biology and Medicine*.
31. De Fauw, J., Keane, P., Tomasev, N., Visentin, D., Driessche, G. V., Johnson, M., . . . Cornebise, J. (2017). Automated analysis of retinal imaging using machine learning techniques for computer vision. *F1000Research*, *5*, 1573. doi:10.12688/f1000research.8996.2
32. Grimm, L. J., Zhang, J., & Mazurowski, M. A. (2015). Computational approach to radiogenomics of breast cancer: Luminal A and luminal B molecular subtypes are associated with imaging features on routine breast MRI extracted using computer vision algorithms. *Journal of Magnetic Resonance Imaging*, *42*(4), 902–907. doi:10.1002/jmri.24879
33. Guerra, C. J., Hackett, M., Carrola, J. T., Couvillion, Warren C., J., Chambers, D. R., & Desai, S. A. (2014). Real-time passive tracking for multi-touch medical modeling and simulation. *Studies in Health Technology and Informatics*, *196*, 142–146.
34. Gupta, A., Ayhan, M. S., & Maida, A. (2013). Natural image bases to represent neuroimaging data. *Proceedings of the 30th International Conference on Machine Learning*, *28*(3), 987–994. Atlanta, Georgia. Retrieved from <http://proceedings.mlr.press/v28/gupta13b.pdf>
35. Hosseini-Asl, E., Gimel'farb, G., & El-Baz, A. (2016). Alzheimer's disease diagnostics by a deeply supervised adaptable 3D convolutional network. arXiv:1607.00556 [cs.LG]. Retrieved from <https://arxiv.org/abs/1607.00556>
36. Payan, A., & Montana, G. (2015). Predicting Alzheimer's disease: A neuroimaging study with 3D convolutional neural networks. arXiv:1502.02506 [cs.CV]. Retrieved from <https://arxiv.org/abs/1502.02506>

37. Plis, S. M., Hjelm, D. R., Salakhutdinov, R., Allen, E. A., Bockholt, H. J., Long, J. D., . . . Calhoun, V. D. (2014). Deep learning for neuroimaging: A validation study. *Frontiers in Neuroscience*, *8*, 229. doi:10.3389/fnins.2014.00229
38. Kim, J., Calhoun, V. D., Shim, E., & Lee, J. H. (2016). Deep neural network with weight sparsity control and pre-training extracts hierarchical features and enhances classification performance: Evidence from whole-brain resting-state functional connectivity patterns of schizophrenia. *NeuroImage*, *124*, 127–146. doi:10.1016/j.neuroimage.2015.05.018
39. Uddin, L. Q., Supekar, K., & Menon, V. (2013). Reconceptualizing functional brain connectivity in autism from a developmental perspective. *Frontiers in Human Neuroscience*, *7*, 458. doi:10.3389/fnhum.2013.00458
40. Colby, J. B., Rudie, J. D., Brown, J. A., Douglas, P. K., Cohen, M. S., & Shehzad, Z. (2012). Insights into multimodal imaging classification of ADHD. *Frontiers in Systems Neuroscience*, *6*. doi:10.3389/fnsys.2012.00059
41. Nielsen, J. A., Zielinski, B. A., Fletcher, P. T., Alexander, A. L., Lange, N., Bigler, E. D., . . . Anderson, J. S. (2013). Multisite functional connectivity MRI classification of autism: ABIDE results. *Frontiers in Human Neuroscience*, *7*, 599. doi:10.3389/fnhum.2013.00599
42. Brown, M. R., Sidhu, G. S., Greiner, R., Asgarian, N., Bastani, M., Silverstone, P. H., . . . Dursun, S. M. (2012). ADHD-200 Global Competition: Diagnosing ADHD using personal characteristic data can outperform resting state fMRI measurements. *Frontiers in Systems Neuroscience*, *6*. doi:10.3389/fnsys.2012.00069
43. Biswal, B. B., Mennes, M., Zuo, X-N., Gohel, S., Kelly, C., Smith, S. M., . . . Milham, M. P. (2010). Toward discovery science of human brain function. *Proceedings of the National Academy of Sciences*, *107*(10), 4734–4739. doi:10.1073/pnas.0911855107
44. Mennes, M., Biswal, B. B., Castellanos, F. X., & Milham, M. P. (2013). Making data sharing work: The FCP/INDI experience. *NeuroImage*, *82*, 683–691. doi:10.1016/j.neuroimage.2012.10.064
45. Poldrack, R. A., Barch, D. M., Mitchell, J. P., Wager, T. D., Wagner, A. D., Devlin, J. T., . . . Milham, M. P. (2013). Toward open sharing of task-based fMRI data: The OpenfMRI project. *Frontiers in Neuroinformatics*, *7*. doi:10.3389/fninf.2013.00012
46. Miller, K. L., Alfaro-Almagro, F., Bangerter, N. K., Thomas, D. L., Yacoub, E., Xu, J., . . . Smith, S. M. (2016). Multimodal population brain imaging in the UK Biobank prospective epidemiological study. *Nature Neuroscience*, *19*(11), 1523–1536. doi:10.1038/nn.4393
47. Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., & Uğurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*, *80*, 62–79. doi:10.1016/j.neuroimage.2013.05.041
48. Di Martino, A., Yan, C-G., Li, Q., Denio, E., Castellanos, F. X., Alaerts, K., . . . Milham, M. P. (2013). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular Psychiatry*, *19*(6), 659–667. doi:10.1038/mp.2013.78

49. Mueller, S. G., Weiner, M. W., Thal, L. J., Petersen, R. C., Jack, C. R., Jagust, W., . . . Beckett, L. (2005). Ways toward an early diagnosis in Alzheimer's disease: The Alzheimer's Disease Neuroimaging Initiative (ADNI). *Alzheimer's & Dementia*, 1(1), 55–66. doi:10.1016/j.jalz.2005.06.003
50. Holmes, A. J., Hollinshead, M. O., O'Keefe, T. M., Petrov, V. I., Fariello, G. R., Wald, L. L., . . . Buckner, R. L. (2015). Brain Genomics Superstruct Project initial data release with structural, functional, and behavioral measures. *Scientific Data*, 2, 150031. doi:10.1038/sdata.2015.31
51. Jernigan, T. L., Brown, T. T., Hagler, D. J., Akshoomoff, N., Bartsch, H., Newman, E., . . . Dale, A. M. (2016). The Pediatric Imaging, Neurocognition, and Genetics (PING) data repository. *NeuroImage*, 124, 1149–1154. doi:10.1016/j.neuroimage.2015.04.057
52. Craddock, R. C., & Bellec, P. (n.d.). Preprocessed Connectomes Project. Retrieved from <http://preprocessed-connectomes-project.org/>
53. Isaksson, A., Wallman, M., Göransson, H., & Gustafsson, M. (2008). Cross-validation and bootstrapping are unreliable in small sample classification. *Pattern Recognition Letters*, 29(14), 1960–1965. doi:10.1016/j.patrec.2008.06.018
54. Rao, R. B., Fung, G., & Rosales, R. (2008). On the dangers of cross-validation. An experimental evaluation. *Proceedings of the 2008 SIAM International Conference on Data Mining*, 588–596. doi:10.1137/1.9781611972788.54
55. Yan, C., Cheung, B., Kelly, C., Colcombe, S., Craddock, R. C., Martino, A. D., . . . Milham, M. P. (2013). A comprehensive assessment of regional variation in the impact of head micromovements on functional connectomics. *NeuroImage*, 76, 183–201. doi:10.1016/j.neuroimage.2013.03.004
56. Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, 84, 320–341. doi:10.1016/j.neuroimage.2013.08.048
57. Fair, D. A., Nigg, J. T., Iyer, S., Bathula, D., Mills, K. L., Dosenbach, N. U., . . . Milham, M. P. (2013). Distinct neural signatures detected for ADHD subtypes after controlling for micro-movements in resting state functional connectivity MRI data. *Frontiers in Systems Neuroscience*, 6. doi:10.3389/fnsys.2012.00080

Conclusion

AI and machine learning represent the future of military systems, networks, and processes. The imminent adoption of human-machine collaboration by DoD components, as well as the desire of chief U.S. competitors (Russia and China) to adopt fully autonomous systems, makes research in this area of utmost importance. The optimization of AI and machine learning for military applications remains a chief priority for DoD. Machine learning and AI were prominent in the 2014 DoD Third Offset Strategy. Instead of driving fully autonomous military systems, machine learning and AI will function as force multipliers for DoD. This strategy promotes an increased level of human-machine interfacing in military systems without undermining the human's authority in key strategic decisions. This vision emphasizes semi-autonomous vehicles, advanced computing, and guidance and control systems that require human involvement to make decisions in battle contexts. This policy stance, termed "assisted human operations," outlines the need for operating systems that use their computational power to give a tactical advantage through lessening human error, while leaving big-picture strategic goals to the human operator.

This report provides an overview of AI as it applies to HDIAC's eight focus areas. While not all-inclusive, the discussions contained herein demonstrate what is currently possible in a variety of fields as well as where the data are lacking. Although defense applications of AI often only focus on combat operations and weapons development, this report illustrates that is not the case at all. Many other areas, from renewable energy forecasts to medicine, play a pivotal role in defense operations, and AI has the potential to enhance those areas, particularly to benefit warfighter performance and enhance operational efficiency. This report can also be thought of as a road map to future research needs. Areas such as medicine, biometrics, and critical infrastructure have more developed AI potential but might benefit from evaluations centered on ethics and emergency management, whereas CBRN defense and WMD are nascent areas of AI research and deserve deeper analysis.

The desired end-state of AI is to perform highly complex computational tasks, under austere conditions, as a seemingly routine task with optimal efficiency and minimal failure (more efficiently than humans in most cases). As with any technology development, there is a certain amount of risk in the early stages and numerous failures and setbacks along the way. At some logical point after repeated benchmarking and successful proofs of concept, decisions to deploy or go to market are made. Part of the challenge is a heavy industry influence to offload human responsibility to AI. Researchers have made incredible leaps forward in AI since the 1950s, but significant work remains. Continued research into how to use AI and a better understanding of how to identify and mitigate risks are especially crucial now as new breakthroughs occur almost every day and so many organizations, including DoD, rely more and more on AI.