

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Artificial intelligence and machine learning: a perspective on integrated systems opportunities and challenges for multi-domain operations

Ravichandran, Ravi, Chong, Chee-Yee, Smith, Robert

Ravi Ravichandran, Chee-Yee Chong, Robert E. Smith, "Artificial intelligence and machine learning: a perspective on integrated systems opportunities and challenges for multi-domain operations," Proc. SPIE 11746, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications III, 1174606 (12 April 2021); doi: 10.1117/12.2587216

SPIE.

Event: SPIE Defense + Commercial Sensing, 2021, Online Only

Artificial Intelligence and Machine Learning: A Perspective on Integrated Systems Opportunities and Challenges for Multi-Domain Operations

Ravi Ravichandran^{*a}, Chee-Yee Chong^b, Robert E. Smith^c

^aBAE Systems Intelligence & Security;

^bIndependent Consultant;

^cUniversity College London

ABSTRACT

This paper provides a perspective on historical background, innovation and applications of Artificial Intelligence (AI) and Machine Learning (ML), data successes and systems challenges, national security interests, and mission opportunities for system problems. AI and ML today are used interchangeably, or together as AI/ML, and are ubiquitous among many industries and applications. The recent explosion, based on a confluence of new ML algorithms, large data sets, and fast and cheap computing, has demonstrated impressive results in classification and regression and used for prediction, and decision-making. Yet, AI/ML today lacks a precise definition, and as a technical discipline, it has grown beyond its origins in computer science. Even though there are impressive feats, primarily of ML, there still is much work needed in order to see the systems benefits of AI, such as perception, reasoning, planning, acting, learning, communicating, and abstraction. Recent national security interests in AI/ML have focused on problems including multi-domain operations (MDO), and this has renewed the focus on a systems view of AI/ML. This paper will address the solutions for systems from an AI/ML perspective and that these solutions will draw from methods in AI and ML, as well as computational methods in control, estimation, communication, and information theory, as in the early days of cybernetics. Along with the focus on developing technology, this paper will also address the challenges of integrating these AI/ML systems for warfare.

Keywords: Artificial Intelligence, machine learning, multi-domain operations, estimation, control, system of systems

1. INTRODUCTION

1.1 Dartmouth Summer Conference

The term “artificial intelligence” (AI) was coined at a Dartmouth Summer Research Project in 1956 [1]. Attendees included Minsky, Shannon, Simon, Newell and McCarthy, who went on to become luminaries in the field. The institutions they represented (MIT, CMU, and Stanford) continue to be thought leaders and centers of excellence for AI. As in any technical area, AI has seen its ebbs and flows, with interest in the large span of subtopics [2]. Central to these topics are capabilities inspired by human intelligence, especially in the early days with symbolic descriptions, such as perception, reasoning, planning, acting, learning, communicating, and abstraction [3]. Core research and development topics include knowledge representation, inference, logic, uncertainty management, neural and belief networks, reinforcement learning, and rule-based systems [4] [5].

1.2 Early Definitions of AI and ML

Since its beginnings, what is meant by “AI” has continuously changed. Following the Dartmouth project, this loosely defined field largely replaced cybernetics, which was a clear outgrowth of work that is now estimation and control. However, largely due to Simon and Newell’s work on general problem solver, symbolic AI emerged as more substantial strand of work. In parallel, so-called neuromorphic AI (which has always only had superficial similarity to brains) was a significant theme, largely through Rosenblatt’s much-hyped work on perceptrons, until that work collapsed under the

*ravi.ravichandran@baesystems.com

weight of Minsky and Papert's book *Perceptrons* [54]. Rules, frames, and similar logical structures came to dominate AI in the 70s and 80s, particularly in defense circles. This work collapsed under the weight of the Lighthill report [55] in the UK and a similar report by the American Study Group in the US, which illustrated that this strand of AI was failing to live up to its hype, as well.

The "AI Winter" of the 1980s (largely precipitated by withdrawal of D/ARPA funds from "AI" projects, and the discovery that knowledge engineering of large systems was not commercially viable in the private sector) led to a re-labelling of rule based expert systems AI as "decision support systems." That is an interesting term, as it implies that human decision-making remained in the driver's seat.

The revival of neuromorphic AI (and, in fact, a range of related ML algorithms) has been precipitated by commercial effects: the availability of cheap and free data on the internet, and the dramatic fall in the price of computation. Thus, when most people say AI now, they almost exclusively mean ML.

An interesting, related effect can be observed by looking at successive years of Gartner Hype Cycle. Many AI topics simply vanish after coming over the top of the curve. This is one of the reasons that AI seems to wax and wane sinusoidally. It is not the same AI from cycle to cycle. It is just that new technologies are labelled AI, just before they fail to meet hype, and then become more useful tools in the systems' engineering toolkit.

This continuous change in what "AI" means partially due to confusion caused by the name itself. AI's precursor term, *cybernetics*, was described by its creator, Wiener [58], as "the scientific study of control and communication in the animal and the machine", and the term itself derives from the Greek for "governor." It is now defined as "a transdisciplinary approach for exploring regulatory systems—their structures, constraints, and possibilities." A variation of this transdisciplinary approach also took root in systems engineering and led to growth in fields such as information theory, estimation, controls, and optimization led by pioneering work by Bellman, Dempster, Narendra, Bryson, Ho, Shannon who led pioneering work in adaptive control, dynamic programming, learning automata, and information theory [42][43][44][45][46].

In contrast, the term "AI" suffers from the two main conflicting meanings of the word "artificial":

- It can have the sense of artificial light. Such light is, in fact, light, but it is created by human-made, rather than natural sources, like the sun.
- It can also have the sense of artificial flowers, which are not flowers, but a functional imitation of them.

The history of AI, and its hype cycles, are partially because of a search for a "magic bullet" that will create AI in the sense of artificial light. That is, a technology that will fundamentally enable machine reasoning that is the same as human reasoning, just man made. As various attempts at this magic bullet have failed to hit their mark, they have been discarded as "not AI." Usually, this has meant they become a tool of mainstream programming and (systems) engineering. Another reason for the lack of an AI "definition" is because of AI's role in the development lifecycle. AI is mostly considered as a "Research Pursuit". Adopted AI is just software, i.e. the "AI effect". For example for a route planning apps like Waze, once adopted as software, do not seem "intelligent" anymore. This is best summarized by Tesler's Theorem "AI is whatever has not been done yet" [61].

Finally, one of the main reasons AI has not hit the mark is because it continues to be a very difficult problem. Fifty years after the Dartmouth conference, McCarthy comments that "AI was harder than we thought" and Minsky that "Easy things are hard" [2].

This leads to the current state of the field, where Artificial Intelligence (AI) and Machine Learning (ML) today are used interchangeably or together as AI/ML and ubiquitous among many industries and applications and where AI largely means ML. Yet what AI means and how near it is to human reasoning remains a subject of debate.

The rest of this paper is structured as follows. Section 2 Innovations and Application of ML; Section 3: Defining AI and ML; Section 4: ML Successes and AI Challenges; Section 5: Representation and Inference; Section 6: National Security Interests; Section 7 AI and ML National Security Successes; Section 8 Multi-Domain Operations: System of Systems; Section 9: Integrated Systems Challenges of AI for MDO; and Section 10: Summary and Conclusions.

2. INNOVATIONS AND APPLICATIONS OF ML

2.1 Since the 2000s

The recent explosion, primarily based on ML, a subset of AI, began in the early 2000s with the explosion of data and the availability of computational resources. In the last decade, this explosion was spurred by the confluence of ML algorithms, specifically deep neural networks, large data sets, and fast and cheap computing, as shown in Figure 1 (adapted from [60]). Algorithms focused on data conditioning, modelling, classification and regression. Work on data sets focused on collection, storage, and retrieval, and computing on CPUs and GPUs. Along with academic institutions, leaders in the field now include Amazon, Facebook, Google, IBM, NVIDIA. Further, with the publicized victories of IBM Watson on Jeopardy, DeepBlue on Chess, and then DeepMind on Go, AI/ML has not only demonstrated impressive technical feats but also captured the public imagination and became a household word.

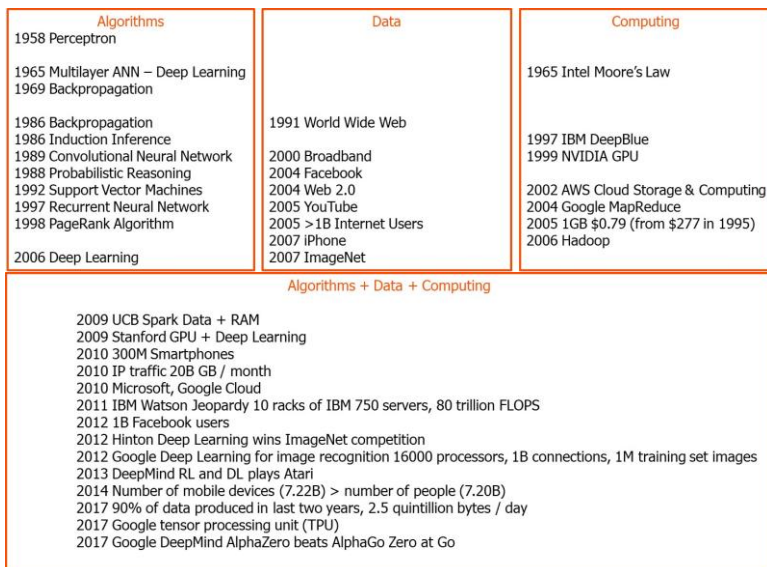


Figure 1. ML and Confluence of Algorithms, Data, and Computing

2.2 Competitive Landscape

The competitive landscape for ML, which used to be niche, now spans a wide range of *applications*, *industries*, and a *complete technology stack*. *Applications* include enterprise intelligence, enterprise functions, autonomous systems, and agents. *Industries* include agriculture, education, finance and investment, legal, logistics, materials, retail, and healthcare. The *technology stack* includes agent enablers, data science, machine learning, natural language processing, data capture, storage, retrieval; open source libraries, and computing hardware [7]. However, note that the combination of Data Science expertise and subject matter expertise (SME) is not easily transferable. For example, expertise in medical data analytics will not easily transfer to geospatial analytics.

2.3 Innovations and Highlights

In the last decade, many innovations and highlights propelled ML. Three of note include Deep Learning, ImageNet, and AlphaGo. Deep Learning (DL) [9] is the pioneering work of LeCun, Bengio, and Hinton, whose insight into representation and abstraction came from expertise in image analysis and speech recognition from three decades of work. Their innovation was to build a common framework for learning multiple levels of representation and abstraction. They also comment on their debt to the availability of large data sets and increase in computational power. Note that this debt is not just to large data, but large sets of *labelled* data. The insight into building large 'labelled' data comes from ImageNet [10] and the work led by Li. Not only did ImageNet build a large, labelled data set, but it also created a knowledge ontology, crowdsourced data set labelling, and metrics to benchmark the data set. Interest in AI and games has always been a topic of interest in the community, but DeepMind's AlphaGo, with neural networks trained by subject matter experts and reinforcement learning by self-play, defeated a world champion in the ancient game of Go. AlphaGoZero [11] took AlphaGo a step further; presented with only rules of the game, it mastered the game by self-play without human knowledge training the system.

2.4 Current ML Maturity

As in any field, maturing in both technology and application there are debates. In ML, these debates include the metrics and reproducibility of results [12]; lack of insight into why some algorithms, for example Support Vector Machines, Principal Component Analysis, work and others do not; and concern that there are no rigorous criteria for choosing one ML or DL architecture over another [13], or standardized benchmarks and metrics [14]. Further, practitioners comment that, as shown in Figure 3, in most applications, the actual work of ML is primarily focused on conditioning and preparing the data [15] or even constructing the data-bases themselves [16], with the actual application of ML algorithms becoming a limited part of making these “automatic” and “general purpose” methods succeed.

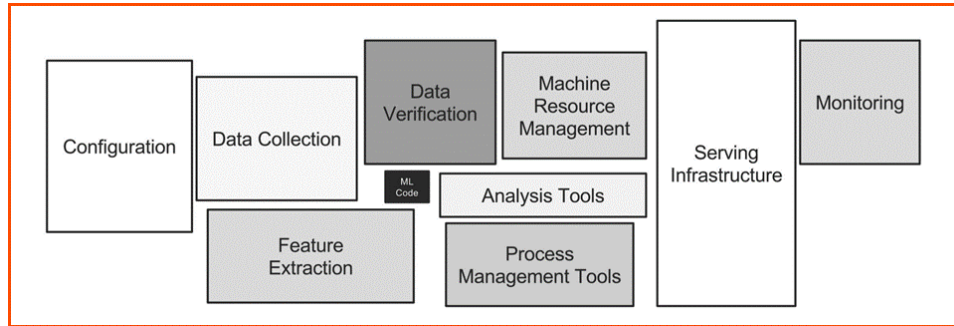


Figure 2. Data processing, configuration and management, and not ML, is the crux implementation

2.5 Computational Tools

With the widespread interest in ML and the applications to data, a practitioner has a vast array of tools across the technology stack. At the forefront are Amazon’s SageMaker [17], Facebook’s PyTorch [18], and Google’s TensorFlow [19] along with many other tools. These tools allow a practitioner to build a ML system for classification and regression. These platforms also offer the practitioner the democratization of ML with entire support ecosystems built around the platforms and have contributed to the acceleration of ML adoption. Another topic is the ML “Marketplace” growth – trained models and training data (to include synthetic data) are now in demand and able to be purchased. This is an entirely new sales channel that has made the ML market more resilient.

The introduction of ML tools as canned algorithms represents a subtle but important shift. In the past, the utilization of ML usually required careful representation design, carefully schematizing the data required to fit a custom-designed application. Now data representations and schemata are more likely to be taken off-the-shelf and pipelined into these canned algorithms. While this makes ML more easily and widely applicable, it also obscures assumptions made in the design of the data schema from the ML practitioner, which can obscure insights on the problem at hand, and it can lead to unseen biases and oversights.

This is part and parcel of the reason that explanations for the results that an ML system produces often remain elusive [20] and the continued exploration about how these systems operate on a theoretical level [21].

Along with software computational tools, advances in hardware, led by pioneers such as NVIDIA for graphics processing units (GPUs), pioneers in cloud computing, and advances in microelectronics and custom “AI chips” have propelled these advances.

3. DEFINING AI AND ML

3.1 ML Definitions

The technical themes of ML, as shown in Figure 2, include clustering (creating discrete groups from unlabelled data), regression (creating weights to match independent data), classification (creating a mapping from data to labels), data completion (inferring missing or withheld data), and model learning (choosing an ideal representation for data) [8]. This has led to a demand from a wide range of industries for ML and specifically for expertise in Data Science. Thus, along with ML, the desired background for data scientists includes proficiency in mathematics, statistics, programming, data processing, data visualization, and application domain expertise [6].

Clustering	Regression	Classification	Data Completion	Model Learning
K-Means	LASSO	SVM	Matric Completion	t-SNE
EM	Ridge Regression	Deep NN	Probabilistic-Graphical Models	Kernel PCA
DBSCAN	Logistic	Naïve Bayes		MDL
Dirichlet Processes	SGD Linear	KNN		Structure Learning

Figure 3. ML for Clustering, Regression, Classification, Data Completion, and Model Learning

3.2 AI Definitions

AI today lacks a precise definition. For example, three recent definitions include

- “Artificial intelligence is an extremely broad discipline, defined in many different ways for many different purposes we use the term to mean a variety of information processing techniques and technologies used to perform a goal-oriented task and the means to reason in the pursuit of that task” (2018 Defense Innovation Board [22]).
- “An interdisciplinary hub for work in computer science, AI, data science, and related fields” (2018 MIT College of Computing announcement [23]).
- “AI has evolved towards a broadly applicable engineering discipline in which algorithms and data are brought together to solve a variety of pattern recognition, learning, and decision-making problems. Increasingly, AI intersects with other engineering and scientific fields and cuts across many disciplines in computing” (2017 “A Berkeley View of Systems Challenges for AI” [24]).

4. ML SUCCESSES AND AI CHALLENGES

4.1 ML Successes and AI Challenges

In recent publications and presentations, leading researchers and practitioners of AI and ML have commented on the impressive feats, primarily of ML, and commented that there still is much work to accomplish in order to see the many benefits implied by the term “AI”. For instance, perception, reasoning, planning, acting, learning, communicating, and abstraction as in real, human intelligence.

4.2 AI today is mostly ML

In “AI – The Revolution Hasn’t Happened Yet” [25] Jordan comments that AI today is mostly what has been called ML. It blends ideas from many disciplines and designs algorithms that process data, make predictions and helps make decisions. However, high-level reasoning and thought remain elusive. Algorithms are assembled ad-hoc and lack an engineering discipline with its principles of analysis and design. (Classical) AI, with a focus on meaning and reasoning into systems that perform natural language processing, that infer and represent causality, that create and use computationally tractable representations of uncertainty; and that create and use systems that formulate and pursue long-term goals, have not made significant strides as a part of the recent surge of interest in AI. The current public dialog on AI focuses on a narrow subset of industries academic efforts, and risks limiting the challenges and opportunities presented by the full scope of AI.

4.3 Focus on Computing rather than Encoding Knowledge

In “The Bitter Lesson” [26], Sutton comments that the advances in AI are based on leveraging computation and not on leveraging domain specific knowledge. For example, in chess and Go, rather than building on an understanding of the structure of the games, the AI systems that bettered human master players did so through search and learning based on fast computational hardware and software applied at very large scales. Likewise, Sutton makes similar observations for computer vision. He summarizes that deeper work in AI begins with building knowledge into agents – in this example, knowledge about the structure of vision. By contrast, in the recent advances for computer vision, the focus is on search and learning. While this helps in the short term, it cannot be relied on to scale in the long term.

For example, consider object detection in vision, where the problem is posed as a hierarchy of detection problems [47]. The solution begins with detecting edge regions, then constructing line and curve segments, then completing and grouping regions and objects, then separating foreground and background, then labelling specific foreground objects and

using them for recognition. At each stage, knowledge is encoded. For example, curve reconstruction in the presence of uncertain information brings in global information such as line and curve continuity when local edge information is either lost or noisy due to insufficient brightness or contrast. Domain knowledge representation is at the focus of methods such as Pattern Theory [48] and Perceptual Organization [49]. Correspondingly, one of the main challenges is finding the appropriate computation tools for inference that match the complex knowledge representations of the problem. Thus, hence, as Sutton observes the focus on computation search is at the expense of these important steps in knowledge representation.

4.4 Associating but not Explaining

In “The Book of Why” (and summarized in the article “To Build Truly Intelligent Machines, Teach Them Cause and Effect”) [27] [28], Pearl comments that the focus of ML is broadly on data association. Today, with large storage and fast computation, ML applications find hidden and non-intuitive associations on very large sets. These implementations are complex and non-trivial. Yet, they have not progressed beyond solving data association problems, using ML techniques that were, in fact, available a generation ago. They are unable to provide an explanation for their results and unable to connect cause and effect. Along with the ability to associate, in order to build the ability to reason or explain, Pearl comments that these systems need to build the ability for causal reasoning, to include association, intervention, and counterfactual reasoning.

5. REPRESENTATION AND INFERENCE: SEMANTIC DESCRIPTION AND NUMERICAL COMPUTATION

5.1 Representation and Inference

Consider ML successes and AI challenges from the vantage point of representation and inference. Progress addressing challenges to methods in perception, reasoning, planning, acting, learning, communicating, and abstraction have been made when the representation of these semantic descriptions have been matched with inference based on numerical computation. This progress has come with tradeoffs. As noted above by Jordan, Sutton, and Pearl, representation based on building domain-based model is hard, and the advances in the access and storage of data and access to fast and cheap computation have focused on statistical models and computation, not these challenges. Primarily “they have limited and narrow function,” which is the observation from LeCun in response to Facebook's progress on speech recognition [30]. These complex tasks are built from a sequence of numerical computations and excel at the specific task, but do not generalize, and practitioners observe that they are easily fooled. For example, auto grading for on-line learning can be fooled by writing a few sentences and a number of key words rather than writing a complete, coherent paragraph [31].

5.2 Semantic Description and Numeric Computation

Even though large and complex, ML progress to date has been made when representation of semantic description by statistical models has been matched by inference based on numerical computation of these statistical models. This approach works well for a family of particularly “data-oriented”, problems, where the inherent semantic descriptions are based on statistical models.

In “Building Machines that Learn and Think Like People” [29] Lake and colleagues also comment that recent advances in ML have focused on discovering high dimensional, abstract features based on statistical pattern recognition with a focus on computation. Lake and colleagues comment on their inspiration to draw from human cognition, and that “human cognitive abilities remain difficult to understand computationally”. This is similar to the observation Sutton makes that contents of the mind are “irreducibly complex,” and the challenge is to find “meta-methods that can find and capture this arbitrary complexity”.

However, they also comment that their focus is on finding methods for these difficult computational problems. These methods include incorporating models from physics and psychology. From physics, the models will both constrain and guide global search, and from psychology, models for punishment and reward will shape cost or utility functions in both spatial and temporal domains. Earlier, we commented on the work from Bayesian Pattern Theory and Perceptual Organization. Foundational principles for both these methods also include models from physics and psychology. Next Lake and colleagues comment on using models from physics and psychology as building blocks to build causal models not only to predict but also to explain observed data.

5.3 Waves of AI

These challenges go back to the early days of AI, and are succinctly described in DARPA Perspectives on AI [32] as the “Three Waves of AI (Figure 4). The 1st wave was built on expert systems based on symbolic handcrafted knowledge – with results that are explainable but system that are not scalable. The current 2nd wave is built on statistical learning – with results that are not explainable but systems that are scalable. They comment that the future 3rd wave, based on contextual adaptation, will focus on both results [59] that are explainable and systems that are scalable, where representation include symbolic information and statistical information and inference includes computation frameworks and methods for symbolic and statistical models. The application of AI to warfare is also captured by DARPA’s work on Mosaic Warfare [63] with a focus on AI and autonomous systems for decision centric operations.




Figure 4. DARPA “Three Waves” of AI

6. NATIONAL SECURITY INTERESTS

Recent United States National Security community documents concerning AI have focused on technology, policy, workforce, and problems of national interest.

6.1 DoD AI Strategy

The primary focus areas of the DoD AI Strategy are: Delivering AI-enabled capabilities that address key missions; Partnering with leading private sector technology companies, academia, and global allies and partners; Cultivating a leading AI workforce; and Leading in military ethics and AI safety. The Joint Artificial Intelligence Center (JAIC) is a focal point of the DoD AI Strategy to accelerate, coordinate, and deliver AI capabilities [33].

6.2 ODNI AIM Strategy

The ODNI (Office of the Director of National Intelligence) AIM (Augmenting Intelligence Using Machines) strategy provides guidance on: Investment Strategy; Policy and Authorities; Workforce Strategy; Industry Partnership Strategy; Roles for USG (United States Government) Agencies, National Labs, FFRDCs (Federally Funded Research and Development Centers), UARCs (University Affiliated Research Centers), Commercial and Academic Institutions; Five Eye Foreign Partner Engagement; AI Assurance – Secure and Maintain Competitive Advantage; Outreach / Communications Strategy; and Governance [34].

6.3 Other Reports

Other recent reports also include MIT/Lincoln Laboratory *Artificial Intelligence: Short History, Present Developments, and Future Outlook* [35], the National Science and Technology Council [36], Congressional Research Service [37], and the Department of Energy [38] and the Government Accountability Office *Artificial Intelligence* [39].

One of the observations in the MIT/Lincoln Laboratory report is the similarities and difference of AI/ML between the consumer sector and national security. Along with the observation related to trust and high-end users for national security, they also make a critical observation regarding the consequence of actions, i.e., AI/ML used for warfare has high consequence of actions, hence also the high priority of ethics in AI for national security [22].

7. AI AND ML NATIONAL SECURITY SUCCESSES

7.1 Mission Applications

National Security applications built on AI/ML span business applications, engineering sub-systems and systems; IT, healthcare, and operations and maintenance. In addition, there is continued investment and development across all these applications.

Sub-system and systems applications include sensor and data fusion, planning, information retrieval and EW. These AI/ML methods rely on explicit knowledge of the problem and include methods such as rule-based expert systems, probabilistic reasoning, and automated planning. The advantages of these methods include explicit representation of knowledge, sound representation of uncertainty, explainable algorithms build trust with users, as well as amenability to standard system engineering, test and evaluation, and V&V. The disadvantages of these methods include handcrafted knowledge (that is difficult even for some simple problems), reasoning requiring complicated algorithms (e.g., inference on Bayesian networks), and development that requires expertise and custom software.

The recent advances in ML bypass the need for explicit knowledge representation and exploit data directly to solve these problems. ML methods used to solve problem directly from data include standard unsupervised learning methods, Deep learning and Reinforcement learning. Successes in building component functions for many applications include object recognition, speech recognition, and language machine translation. The advantages of these methods include data available for training many algorithms, no need for sophisticated reasoning not, training only on data, and the availability of open source deep neural network software and convenient development environments. The disadvantages include algorithm sensitivity to training data, the potential for training data bias, algorithms that are hard to explain and difficult to trust, difficulty in system integration (due in large part to lack of explainability), and that competitors and adversaries have access to these tools, advancing computing power speed up for training and execution (for example, through neural network chips).

7.2 AI and ML Successes

DARPA is at the forefront of developing AI/ML for national security applications. Recent examples include applications to RF. Air Combat, C2, and system test and evaluation, validation and verification (TEVV).

The radio frequency (RF) domain is becoming challenging to operate in as the spectrum becomes ever increasingly crowded and the systems that use it become increasingly complex, agile, and less distinct from one another. Recent work using deep learning has pushed the bounds two orders of magnitude with >95% classification accuracy at populations of 10,000 devices [63]. Transferring learning provides one means to alleviate the need for operational data, enabling learning under similar tasks or from simulated data, with RF transfer learning demonstrating 5% improvement due to generalization in some cases and only 2% degradation in constrained cases over training on unavailable operational data [64]. Not all emitters are known ahead of time, however, so the approach needs to extend to classify known emitters as well as detect novel, operational devices while keeping false positive rates well below 1% [65]. Additionally, labels for all devices may not be provided, but downstream processing in the system might require signals separated by emitter, so the solution needs to further handle unsupervised clustering of emitters [66].

Air combat evolution (ACE) aims to test an AI fighter pilot for autonomous air to air engagements. This program made headlines when the AI beat a human pilot 5 times [67]. Other examples include learning algorithms to assess and create imbalance in complex games [68], courses of action (COA) for force-on-force engagements. Objective is for AI COAs to help red teams develop creative adversaries for blue force training and tactics development.

In addition, TEVV is a serious challenge for AI/ML systems. Challenges include the lack of a well-defined performance space and the lack of predictability and transparency in AI/ML systems, especially the highly nonlinear behaviour of deep learning systems. The Assured Autonomy program addressed the challenges of V&V of learning-enabled autonomous systems [69].

DARPA is also exploring AI/ML programs for mission-integrated network control mission goals inform how communications are routed in the network, Engineering AI systems to assemble complex systems involving many AI components, O-Drive for automated, cross functional design of systems of systems like planes, car, and ships.

8. MULTI-DOMAIN OPERATIONS: OPPORTUNITIES AND CHALLENGES

8.1 Joint Warfighting Concept

Along with Joint Logistics, Long Range Fires, and Information Advantage, Joint All Domain Command and Control or Multi-Domain Operations (MDO) is one of the tenets of the Joint Warfighting Concept. As shown in Figure 5, an example of MDO is an “Any Sensor to Any Shooter” “Kill Web” for “Time Sensitive Targeting”, an ensemble of platforms, sensors, weapons, and communications integrated to provide decision advantage at machine speed [52][53].

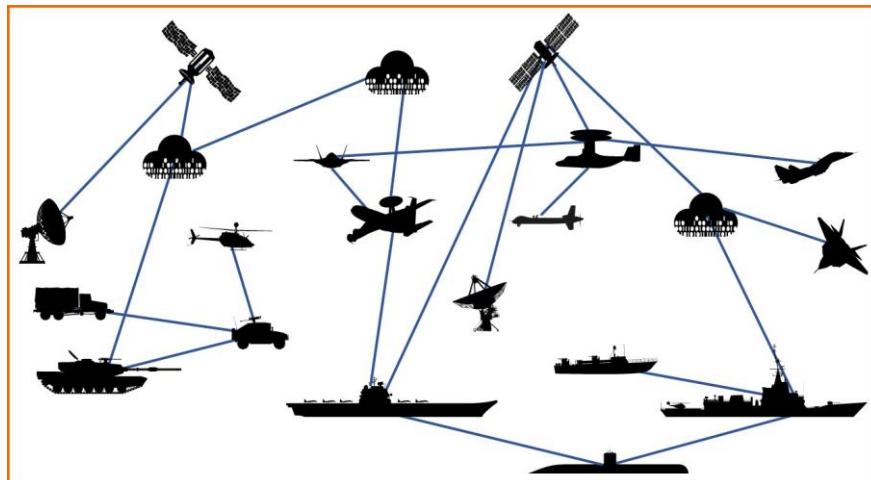


Figure 5. MDO Any Sensor to Any Shooter Kill Web for Time Sensitive Targeting

8.2 MDO Operationalizing AI and ML

There are many technical challenges that need to be addressed in order to provide decision advantage at machine speed. These include a focus on data science, but also on integrated systems that include sub-systems, systems, and system-of-systems. One of the integrated systems challenges is the Intelligence, Surveillance, and Reconnaissance (ISR), Command Control (C2), and Battle Management (BM) mission loop (Figure 6). For MDO, these are addressed both from a system and system-of-systems perspectives.




Figure 6. MDO Sub-Systems, System or Systems-Systems ISR, C2, BM

From a system perspective, ISR, C2, and BM each have their challenges that include data, synchronization, and effects. These systems consist of infrastructure that typically includes platforms and sensors, communications and data links, databases and data management, and visualization and dissemination technology. These systems are built to help the decision maker collect data, comprehend information, anticipate patterns and courses of action, and disseminate knowledge for a variety of missions. Those missions can be characterized by their timelines. Targeting and close air support have timelines in the order of seconds; tipping and cueing sensors and situational awareness in the order of minutes; force protection and special operations in the order of hours; and trend analysis and intelligence preparation in the order of days.

Along with infrastructure, there are a suite of analytic tools that also support ISR, C2, and BM. Analytic tools that support collection include tools for command control that coordinate and synchronize assets across multiple domains and organizations; tools for mission planning that plan routes for multiple aircraft simultaneously and in real-time; and tools for sensor resource management that balance multiple, competing objectives while simultaneously tasking multiple sensors. Analytic tools that support comprehension include single-intelligence target tracking that filters noisy data and produces tracks based on sensor data; and multiple-intelligence correlation and fusion that takes data from two or more sensors and produces a consolidated track picture. Analytic tools that support anticipation include tools for pattern discovery that learn ordinal and temporal patterns and build environmental and normalcy models; rule-based alerts that trigger based on partial and approximant matches; and behavior characterization that establishes causality, predicts intent, explains anomalies, and analyzes alternatives. This description is similar to the challenges presented in the recent NSCAI report [3], in Section 3 “AI and Warfare”, specifically describing “Ways to Operationalize AI”, that are system of system operations (Figure 7).




Figure 7. NSCAI Ways to Operationalize AI

These analytic tools support a wide variety of missions, and as such, can be hosted and accessed on a variety of platforms, on ground stations, or even at remote sites. MDO will only exacerbate these challenges. One of the challenges is that these ISR, C2, and BM systems are mission-specific and existing infrastructure was designed to meet specific mission requirements and not designed to be shared. But, for MDO, the objective is to share data, across services and even joint coalitions. Initiatives across the national security communities for enterprise and shared services include the JIE (Joint Information Environment), ICITE (Intelligence Community Information Technology Enterprise), and Defense Intelligence Information Enterprise (DI2E). Challenges that these projects are working through include updating IT infrastructure, developing a framework for shared services, and enabling cloud-based solutions.

Data, in its varied shapes and forms, is a fundamental thread through ISR, C2, and BM. Challenges include managing (1) quantity— it has been reported for more than a decade that the national security community uses only a fraction of the quantity of data it collects. (2) Rate— any multi-intelligence mission experiences the challenges of the rate with data flowing in real-time, or in batches, or as a stream that is required to stay on for weeks or months. (3) Variety – MDO is rich in variety, including geospatial intelligence, human intelligence, measurement and signature intelligence, signal intelligence, and open-source intelligence, sources that deliver structured, unstructured, semi-structured, and dynamic data. (4) Truthfulness – it is critical to establish truthfulness to the intelligence that is produced to ascribe trustworthiness to the data. (5) Mission value – finally, provide mission value to the decision maker.

For example, consider the vignette shown in Figure 8 that includes an ensemble of platforms, sensors, weapons, and communications on a strike mission. The vignette begins with:

1. WAS (wide area surveillance) platform that tracks multiple moving targets.
2. Then, mission commander assigns FMV (full motion video) platform to PID (positively identify) a track and designate the track as a target.
3. To confirm the target, WAS captures a SAR (synthetic aperture radar) image of the designated target, refining its coordinates.
4. Then WTP (Weapon Target Pairing) algorithm matches available shooters to task against the target.
5. BMC2 operator then selects the WTP options and transmits strike to task.
6. Re-tasks the platforms to new missions.

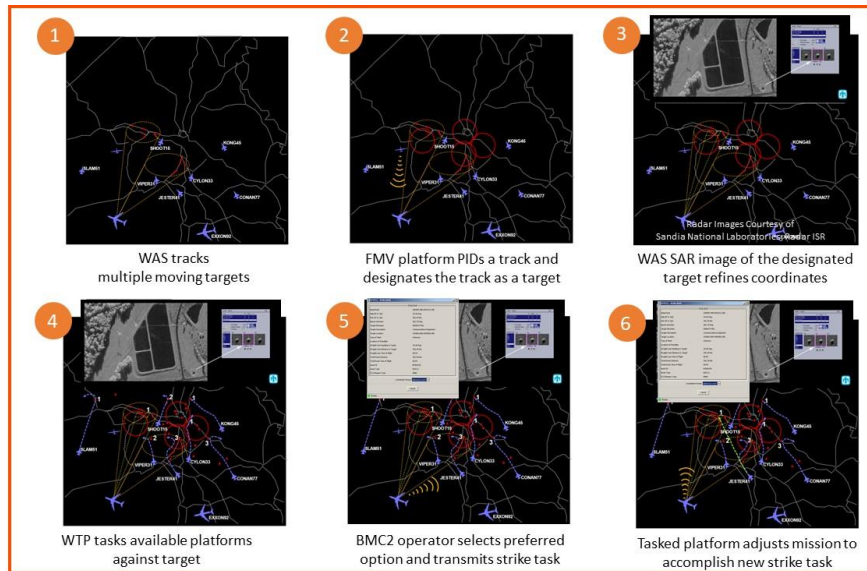


Figure 8. MDO ISR, C2, BM Vignette

8.3 MDO System of Systems

The DoD has addressed the challenges of System of Systems in many programs. These programs include new initiatives, strategic programs, and acquisition programs across both DoD and IC. For System of Systems, challenges include (1) management and oversight, (2) operational environment, (3) implementation, and (4) engineering design and considerations as summarized in Figure 9 (from Table 2-1 in the DoD Systems Engineering Guide for Systems of Systems) [70]. From an engineering perspective one of the primary SoS challenges is integrating systems that were not initially designed to be a part of a SoS, thus these systems objectives many not align with SoS objectives.

Environment	System	System of Systems
Management and Oversight		
Stakeholder Involvement	Clearer set of stakeholders	Stakeholders at both system level and SoS levels (including the system owners), with competing interests and priorities; in some cases, the system stakeholder has no vested interest in the SoS; all stakeholders may not be recognized
Governance	Aligned PM and funding	Added levels of complexity due to management and funding for both the SoS and individual systems; SoS does not have authority over all the systems
Operational Environment		
Operational Focus	Designed and developed to meet operational objectives	Called upon to meet a set of operational objectives using systems whose objectives may or may not align with the SoS objectives
Implementation		
Acquisition	Aligned to ACAT Milestones, documented requirements, SE with a Systems Engineering Plan (SEP)	Added complexity due to multiple system lifecycles across acquisition programs, involving legacy systems, systems under development, new developments, and technology insertion; Typically have stated capability objectives upfront which may need to be translated into formal requirements
Test & Evaluation	Test and evaluation of the system is generally possible	Testing is more challenging due to the difficulty of synchronizing across multiple systems' life cycles; given the complexity of all the moving parts and potential for unintended consequences
Engineering Design & Considerations		
Boundaries and Interfaces	Focuses on boundaries and interfaces for the single system	Focus on identifying the systems that contribute to the SoS objectives and enabling the flow of data, control and functionality across the SoS while balancing needs of the systems
Performance & Behavior	Performance of the system to meet specified objectives	Performance across the SoS that satisfies SoS user capability needs while balancing needs of the systems

Figure 9. Comparing Systems and Systems of Systems

8.4 Evolution of MDO System of Systems

The solutions to these MDO integrated systems challenges will have a systems view, be time critical, and will support validation and verification. The systems view will address tightly coupled system of systems with large amounts of data and bi-directional information flowing through multiple levels of security. Time criticality will address many time-varying decision loops from milliseconds to minutes and hours and varying degrees of data content from rich to sparse. Validation and verification will build trust through explainable decisions, traceable back through each decision point, and support high consequence actions.

As observed in Section 1.2, pioneers understood the close connection between AI/ML and estimation and control for systems challenges. Even four decades ago, as shown in Figure 10, Saridis [50] envisioned an autonomous system to include estimation, control, and learning; that design begins with a-priori knowledge of underlying models for estimation and control. In the MDO case, due to near-peer attacks and information manipulation the system will change due to dynamics of the environment, and thus underlying models will have to be updated.

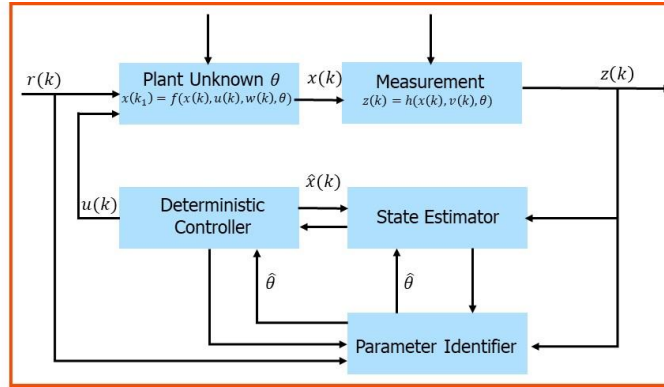


Figure 10. Integrated Systems Estimation, Control, and AI/ML

With advances and the close interconnections between estimation, control, and learning, the view of systems has also evolved. As shown in Figure 10, feedback control methods can tolerate significant model uncertainties and can meet complex system specifications around a specific operating point or range. When there is a significant increase in the operating range due to increasingly complex system dynamics and uncertainties in the environment, techniques to design a class of adaptive controllers that learn optimal control solutions without knowing the full system dynamics or the environment are needed [41]. These will include the integration of estimation and control with AI/ML. That is, there is an ‘inner loop’ for estimation and control, and an ‘outer loop’ for adaptation.

9. INTEGRATED SYSTEMS CHALLENGES OF AI FOR MDO

9.1 MDO Challenges

MDO is in response to near-peer, air-land and air-sea great powers conflicts. MDO system of systems is a highly complex environment made of both legacy and new platform, sensors, communication, weapons, and analytics focusses on heterogeneous and distributed asynchronous operations. ISR will include novel full spectrum sensors, C2 will include complex planning problems, and BM account for peer adversarial threats

In this high intensity conflict against opposing system-of-systems, military operations will still be soldier-centric and the goal of AI enabled systems should not be to replace the soldiers, but to give them tools in their arsenal that improve their survivability and mission effectiveness.




Figure 11. MDO Integrated Systems Estimation, Control, AI/ML and Game Theory

As shown in Figure 11, processes of estimation and control each contribute essential elements to the battlefield. Further, these systems must operate in dynamic battlefields against peer adversaries. Even the most capable systems, which could be augmented with AI/ML, must also integrate learning with estimation and control, so that they can account for multiple environments and changes within environments, as well as to changes to themselves. Finally, Game Theory, the ability for Blue-Red war gaming and the dynamic need to adapt and learn quickly and for ‘counter-AI’ must be an essential capability of these integrated systems.

9.2 Opportunities for Using AI/ML for MDO

Each system in MDO converts observation and state information into decision-making actions and role coordination. This closed loop estimation, control, and AI/ML system, Figure 12, begins with off line models and parameters, composed using a-priori knowledge. In-situ learning monitors the overall performance against mission objectives, adapts according to any mission updates from the operator, and provides assessment for situational awareness and guidance for decision-making. Inputs include sensor measurements, video, target information, communication from other systems, etc. Outputs include target tracks, classifications, weapon commands, communication to other systems) to perform a function (tracking, ID, situation assessment, path planning, weapon management, etc.). When there is a significant change, the closed loop estimation, control, AI/ML system updates the necessary models and parameters to fulfil the mission.




Figure 12. Closed Loop Estimation, Control, and AI/ML

Figure 13 illustrates an example. Analysts have divided a search area into quadrants. Their task is to follow (red) high value targets amidst many (black) non-combatant targets. Based on a-priori information on high value target behaviour, for this mission, Figure 13 left, activity is explicitly constrained to the bottom left quadrant. Thus, target recognition, tracking, path planning, and weapon management for ISR, C2, and BM, are all focussed on the bottom left quadrant and coordinated between a wide area search platform and a narrow area ID platform. As related to Figure 12, these are primarily composed of “inner loop” estimation and control. These models and parameters are built off line a-priori based on prior mission understanding. Also per Figure 12, the “AI/ML learning” outer loop monitors this activity.

When changes in the dynamic battlefield move the area of interest from the bottom left to top right quadrant, it is envisioned that the closed loop estimation, control, and AI/ML system, will be able to identify that the high value targets move to the top right quadrant (Figure 13 right). For MDO, information about change can also come from learned models across domains – a ground based AI-surveillance system in one quadrant could exchange learned target behaviour models with airborne surveillance systems. Further, during conflict or in countries where the government does not hold a monopoly on force, it can be difficult to determine which force, militia, tribe, or other group controls an area [71]. Pictures and videos from social media – while sometimes difficult to verify – can provide information into shifting dynamics on the ground [72]. The coordination between the inner loop and the outer loop is a key design feature, perhaps best explained by Bellman [43] differentiating between a stochastic control process and an adaptive processes based on the information (and uncertainty) related to the system to be controlled. This results in the AI/ML augmenting the estimation and control models and parameters that move the surveillance platforms to the top right quadrant and continue their mission to follow the high value targets.



Figure 13. (Left) Estimation and Control, based on off-line models and parameters prosecutes High Value Targets
(Right) AI/ML Augments Models and Parameters Enabling Estimation and Control to continue mission

9.3 Challenges of Using AI/ML for System of Systems MDO

The academic and national security communities acknowledge that MDO integrated systems will build on the individual advances made in estimation, control, and AI/ML, and that the confluence and integration of these technologies is still a nascent capability, in both theory and practice. A recent DARPA comment succinctly summarizes AI for MDO warfare “Warriors will have to understand what it’s like to have an AI as a trusted team member ...currently, AI isn’t yet ready for prime time, It’s still fragile, opaque, biased and not robust enough, which means it does not yet have trustworthiness” [51].

The design of AI/ML system of systems for MDO will include challenges for (1) system engineering and design (Figure 14), (2) component system development (Figure 15), and (3) integrated system development (Figure 16).

Teaming between warfighter and AI/ML systems	Using AI/ML versus estimation and control	Using data-based ML versus model-based AI
Identifying strengths and limitations of warfighter and systems	Identifying requirements of MDO functions	Availability of models and reasoning approach - AI
Seamless transition between warfighter and systems	Trade-off between using AI/ML versus estimation/control	Models are too hard for reasoning and data are available – ML
Retaining ultimate control of warfighter	Human like intelligent behavior versus mathematical algorithm	Combining data-based ML with model-based AI
	Reasoning and learning versus inference/optimization given models	

Figure 14. System Engineering and Design

Robust, trusted, and high-performance AI/ML components	AI/ML in adversarial environments	AI/ML on the edge with limited communication and computation
Combining AI/ML approaches to improve performance	Lack of real training data; simulated data sensitive to assumptions	Limited communication requires local computation
Performance characterization	Real-time adaptation to adversary behavior	Validating new behaviors after learning in the field
Addressing explainability and interpretability	Protection against adversarial attacks	Limited computation due to size weight and power (SWaP) constraints
Developing trust by testing, explainability, etc	Hardening AI for the field	
Bias in algorithms	Cybersecurity	
Dependence on training data		

Figure 15. Component System Development

AI/ML system test and evaluation, validation and verification (TEVV)	TEVV in national security	Update in the field
Common framework for TEVV	Sensitivity to adversarial behaviour	Updating AI/ML models with limited communication and computing
Confirming components are performing as intended	Slow and sequential process	Validating new behaviors after learning in the field
Metrics for functional performance, robustness, resilience, and trust	Adapting AI/ML to legacy equipment and interfacing with other components	Verifying the applicability of learned models in a new domain to permit sharing of information across domains
Dependence of test results on number of test case		

Figure 16. Integrated System Development

10. SUMMARY AND CONCLUSIONS

This paper provided a perspective on historical background, innovation and applications of Artificial Intelligence (AI) and Machine Learning (ML), including successes, systems challenges, national security interests, and mission opportunities for system problems.

The terms AI and ML today are used interchangeably or together as AI/ML, and ubiquitous among many industries and applications. The recent explosion, based on a confluence of ML algorithms, large data sets, and fast and cheap computing, has demonstrated impressive results in classification and regression and specifically for prediction, and decision-making.

Yet, AI/ML today still lacks a precise definition, and as a technical discipline, it has grown beyond its origins in Computer Science. Even though there are impressive feats, primarily of ML there still is much work needed to accomplish to see the systems benefits of AI such as perception, reasoning, planning, acting, learning, communicating, and abstraction. Recent National Security interests in AI/ML have focused on problems including multi-domain operations (MDO). This has renewed the focus on a systems view of AI.

For the MDO, similar to DARPA's "Mosaic Warfare" and NSCAI's "Operationalizing AI", the problem is posed as a challenge in building and deploying system of systems. Then observed that the opportunities and challenges of AI and ML for integrated systems will draw from both symbolic computational methods in AI and ML and numeric computational methods in control, estimation, communication, and information theory, as in the early days of cybernetics. The confluence and integration of these technologies is still a nascent capability, in both theory and practice. Technical challenges include architecture and system integration; computation; communication, cyber, data and representation, uncertainty management, information and decision-making, validation and verification, and adversarial modelling. The design of AI/ML system of systems for MDO will include challenges for system engineering and design (Figure 14), component system development (Figure 15), and integrated system development (Figure 16).

11. ACKNOWLEDGEMENTS

The authors would like to thank Shaun Donaldson, Steve Homeyer, Steve Jameson, Scott Kuzdeba, Marco Pravia, and Don Widener for their review and comments. This document is not export controlled per IS 2021-1600.

REFERENCES

- [1] J. McCarthy et al, "A Proposal for the Dartmouth Summer Research Project in Artificial Intelligence 1955," reprinted in AI Magazine, vol. 27, no. 4, pp. 12-14, 2006.
- [2] M. Mitchell, Artificial Intelligence A Guide for Thinking Humans, FSG, 2019.
- [3] National Security Commission on Artificial Intelligence Final Report March 2021 <https://www.nsc.ai.gov/>.
- [4] S. Russell and P. Norvig, Artificial Intelligence A Modern Approach, Prentice Hall, 4th Edition 2020.
- [5] P. H. Winston, Artificial Intelligence, Addison Wesley, 2nd Edition 1994.
- [6] S. Page, The Model Thinker, Basic Books, 2018.
- [7] S. Zilis and C. James, "The Competitive Landscape for Machine Intelligence," Harvard Business Review, <https://hbr.org/2016/11/the-competitive-landscape-for-machine-intelligence> November 2016.
- [8] G. Cybenko and P. Chin, "Machine Learning for Multi-Int Fusion" AFRL Multi-Int ID Fusion Workshop, Dayton OH, 2017
- [9] Y. LeCun, Y. Bengio and G. Hinton, "Deep Learning," Nature, May 2015.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li and L.-L. Fi, "ImageNet: A Large-Scale Hierarchical Image Database," in CVPR, 2009.
- [11] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. La, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel and D. Hassabis "Mastering the game of Go without human knowledge," Nature, May 2017.
- [12] M. Hutson, "Artificial intelligence faces reproducibility crisis," Science, February 2018.
- [13] D. Sculley, J. Snoek, A. Rahimi and A. Wiltschko, "Winners Curse? On Pace, Progress, and Empirical Rigor," in International Conference on Learning Representations, 2018.

- [14] D. Blalock, J. J. G. Ortiz, J. Frankle and J. Guttag, "What is the State of Neural Network Pruning?" in Proceedings of Machine Learning and Systems, 2020.
- [15] D. Sculley et al, "Hidden Technical Debt in Machine Learning Systems," in NIPS 2015
- [16] Northeastern University Genesys Lab, "Machine Learning Datasets for RF," <http://genesys-lab.org/mldatasets>, 2020.
- [17] Amazon SageMaker, [Online]. Available: <https://aws.amazon.com/sagemaker/>. [Accessed 19 August 2020].
- [18] PyTorch, [Online]. Available: <https://pytorch.org/>. [Accessed 19 August 2020].
- [19] TensorFlow, [Online]. Available: <https://www.tensorflow.org/>. [Accessed 20 August 2020].
- [20] P. Voosen, "The AI detectives," Science, July 2017.
- [21] National Science Foundation, "Mathematical and Scientific Foundations of Deep Learning" https://www.nsf.gov/news/special_reports/announcements/082520.jsp, August 2020.
- [22] D. I. Board, "AI Principles: Recommendations on the Ethical Use of AI by the Department of Defense," 2018, https://media.defense.gov/2019/Oct/31/2002204459/-1/-1/0/DIB_AI_PRINCIPLES_SUPPORTING_DOCUMENT.PDF.
- [23] MIT reshapes itself to shape the future, <http://news.mit.edu/2018/mit-reshapes-itself-stephen-schwarzman-college-of-computing-1015>.
- [24] "A Berkeley View of Systems Challenges for AI," <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-159.pdf>, 2017.
- [25] "AI – The Revolution Hasn't Happened Yet," 2018. [Online]. Available: <https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7>. [Accessed 19 August 2020].
- [26] "The Bitter Lesson," 2019. [Online]. Available: <http://incompleteideas.net/IncIdeas/BitterLesson.html>. [Accessed 19 August 2020].
- [27] J. Pearl and D. Mackenzie, The Book of Why, Basic Books, 2018.
- [28] "To Build Truly Intelligent Machines, Teach Them Cause and Effect," 2018. [Online]. Available: <https://www.quantamagazine.org/to-build-truly-intelligent-machines-teach-them-cause-and-effect-20180515/>. [Accessed 19 August 2020].
- [29] B. Lake, T. Ullman, J. Tenenbaum and S. Gershman, "Building Machines that Learn and Think Like People," Behavioral and Brain Sciences, 2017.
- [30] WSJ, "Why Some Hate Speech Continues to Elude Facebook's AI Machinery," 9 July 2020, <https://www.wsj.com/articles/facebooks-artificial-intelligence-doesnt-eliminate-objectionable-content-report-finds-11594287000?mod=searchresults&page=1&pos=1>.
- [31] VERGE, "These students figured out their tests were graded by AI — and the easy way to cheat," 2 September 2020, <https://www.theverge.com/2020/9/2/21419012/edgenuity-online-class-ai-grading-keyword-mashing-students-school-cheating-algorithm-glitch>.
- [32] J. Launchberry, "A DARPA Perspective on Artificial Intelligence," 2016. [Online]. Available: <https://www.darpa.mil/attachments/AIFull.pdf>. [Accessed 19 August 2020]
- [33] "SUMMARY-OF-DOD-AI-STRATEGY," <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/SUMMARY-OF-DOD-AI-STRATEGY.PDF>, 2018.
- [34] "The AIM Initiative," <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf>, 2019.
- [35] "Artificial Intelligence: Short History, Present Developments, and Future Outlook," <https://www.ll.mit.edu/media/9526>, 2019.
- [36] "National Artificial Intelligence Research and Development Strategic Plan," <https://www.whitehouse.gov/wp-content/uploads/2019/06/National-AI-Research-and-Development-Strategic-Plan-2019-Update-June-2019.pdf>, 2019.
- [37] "Artificial Intelligence and National Security," <https://crsreports.congress.gov/product/pdf/R/R45178/3>, 2018.
- [38] "Scientific Machine Learning Core Technologies for Artificial Intelligence," <https://www.osti.gov/biblio/1478744>.
- [39] "Artificial Intelligence," <https://www.gao.gov/products/gao-18-142sp>, 2018.
- [40] "JAIC," [Online]. Available: <https://www.ai.mil/index.html>. [Accessed August 19 2020].
- [41] K. Vamvoudakis et al, "Autonomy and Machine Intelligence in Complex Systems," in American Control Conference, 2015.
- [42] R. Bellman, Dynamic Programming, Princeton University Press 1957
- [43] R. Bellman, Adaptive Control Processes, Princeton University Press, 1961
- [44] K. Narendra and M. A. L Thathachar, Learning Automata, Dover Publications, 2012
- [45] A. Bryson and Y-C Ho, Applied Optimal Control, CRC Press, 1975
- [46] R.G. Gallager, Information Theory and Reliable Communication, Wiley, 1968

- [47] S. Casadei and S. Mitter, "Hierarchical Image Segmentation", *International Journal of Computer Vision* 27(1) 1998
- [48] U. Grenander and M. Miller, *Pattern Theory: From Representation to Inference*, Oxford University Press, 2007
- [49] K. Boyer and S. Sarkar, *Perceptual Organization for Artificial Vision Systems*, Springer, 2000
- [50] G. N. Saridis, *Self-Organizing Control of Stochastic Systems*, Marcel Dekker, 1977
- [51] DoD News "Artificial Intelligence is a Work in Progress", January 2021
<https://www.defense.gov/Explore/News/Article/Article/2480288/artificial-intelligence-is-a-work-in-progress-official-says>
- [52] D. Deptula, H. Penney, L. Stuzreim, and M. Gunzinger, "Restoring America's Military Competitiveness: Mosaic Warfare" The Mitchell Institute for Aerospace Studies, 2019
- [53] R. Work, A Joint Warfighting Concept for Systems Warfare, Center for New American Studies, December 2020
- [54] M. Minsky, and S. Papert, *Perceptrons: an introduction to computational geometry*. Cambridge, Mass: MIT Press, 1988.
- [55] J. Lighthill. "[Artificial Intelligence: A General Survey](#)" in *Artificial Intelligence: a paper symposium*, Science Research Council. 1973
- [56] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.
- [57] W. J. Clancey. *Situated cognition: On human knowledge and computer representation*. Cambridge: Cambridge University Press, 1997.
- [58] W. Norbert. *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, Massachusetts: MIT Press, 1948
- [59] G. Marcus and E. Davis, *Rebooting AI Building AI We Can Trust*, Pantheon 2019
- [60] McKinsey & Company "An Executives Guide to AI" <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/an-executives-guide-to-ai>
- [61] [Tesler's theorem and the problem of defining AI - ThinkAutomation](#)
- [62] CSBA "Mosaic Warfare" 2020 https://csbaonline.org/uploads/documents/Mosaic_Warfare.pdf
- [63] J. Robinson and S. Kuzdeba, "Riftnet Radio frequency classification for large populations," IEEE Consumer Communications and Networking Conference (CCNC), 2021.
- [64] S. Kuzdeba, J. Robinson, and J. Carmack, "Transfer learning with radio frequency signals," IEEE Consumer Communications and Networking Conference (CCNC), 2021.
- [65] J. Robinson and S. Kuzdeba, "Novel device detection using RF fingerprints," IEEE Computing and Communication Workshop and Conference (CCWC), 2021.
- [66] J. Stankowicz and S. Kuzdeba, "Unsupervised Emitter Clustering through Deep Manifold Learning," IEEE Computing and Communication Workshop and Conference (CCWC), 2021.
- [67] Air Force Magazine "Artificial Intelligence Easily Beats Human Fighter Pilot in DARPA Trial", Aug 2020
[Artificial Intelligence Easily Beats Human Fighter Pilot in DARPA Trial - Air Force Magazine](#)
- [68] Avila AM, Fonoberova M, Hespanha JP, Mezic I, Clymer D, Goldstein J, Pravia MA, and Javorssek D, "Game Balancing using Koopman-based Learning," invited manuscript accepted to the American Controls Conference 2021
- [69] Neema, S "Assured Autonomy" August 2017, DARPA Industry Day
https://www.darpa.mil/attachments/AssuredAutonomyProposersDay_Program%20Brief.pdf
- [70] DoD Systems Engineering Guide for Systems of Systems, 2008 (accessible on-line)
- [71] Widener, D. V., Mazzuchi, T. A., & Sarkani, S. "Simplifying humanitarian assistance/disaster relief analytic models using activity-based intelligence". *Disaster Prevention and Management*, 2018
- [72] Lewis, K., Dykstra, M. and Widener, D. (2016), *Trend Analysis: The Impact of Syrian Refugees*, BAE Systems, McLean, VA