

SOAR

STATE-OF-THE-ART REPORT (SOAR)
SEPTEMBER 2022



ARTIFICIAL INTELLIGENCE (AI) FOR WEAPONS SYSTEMS

DSIAC-BCO-2022-216

By Sam Chakour
Contract Number: FA8075-21-D-0001
Published By: DSIAC



DSIAC

DISTRIBUTION STATEMENT A
Approved for public release: distribution unlimited.

This Page Intentionally Left Blank

SOAR

STATE-OF-THE-ART REPORT (SOAR)
SEPTEMBER 2022

**ARTIFICIAL
INTELLIGENCE (AI)
FOR WEAPONS SYSTEMS**

SAM CHAKOUR

ABOUT DSIAC

The Defense Systems Information Analysis Center (DSIAC) is a U.S. Department of Defense (DoD) IAC sponsored by the Defense Technical Information Center (DTIC). DSIAC is operated by SURVICE Engineering Company under contract FA8075-21-D-0001 and is one of the three next-generation IACs transforming the DoD IAC program: DSIAC, Homeland Defense & Security Information Analysis Center (HDIAC), and Cybersecurity and Information Systems Information Analysis Center (CSIAC).

DSIAC serves as the U.S. national clearinghouse for worldwide scientific and technical information in 10 technical focus areas: weapons systems; survivability and vulnerability; reliability, maintainability, quality, supportability, interoperability (RMQSI); advanced materials; military sensing; energetics; autonomous systems; directed energy; non-lethal weapons; and command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR). As such, DSIAC collects, analyzes, synthesizes, and disseminates related technical information and data for each of these focus areas. These efforts facilitate a collaboration between scientists and engineers in the defense systems community while promoting improved productivity by fully leveraging this same community's respective knowledge base. DSIAC also uses information obtained to generate scientific and technical products, including databases, technology assessments, training materials, and various technical reports.

State-of-the-art reports (SOARs)—one of DSIAC's information products—provide in-depth analysis of current technologies, evaluate and synthesize the latest technical information available, and provide a comprehensive assessment of technologies related to DSIAC's technical focus areas. Specific topic areas are established from collaboration with the greater defense systems community and vetted with DTIC to ensure the value-added contributions to Warfighter needs.

DSIAC's mailing address:

DSIAC
4695 Millennium Drive
Belcamp, MD 21017-1505
Telephone: (443) 360-4600

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>				
1. REPORT DATE September 2022		2. REPORT TYPE State-of-the-Art Report		3. DATES COVERED
4. TITLE AND SUBTITLE Artificial Intelligence (AI) for Weapons Systems			5a. CONTRACT NUMBER FA8075-14-D-0001	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Sam Chakour			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AND ADDRESS(ES) Defense Systems Information Analysis Center (DSIAC) SURVICE Engineering Company 4695 Millennium Drive Belcamp, MD 21017-1505			8. PERFORMING ORGANIZATION REPORT NUMBER DSIAC-BCO-2022-216	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Technical Information Center (DTIC) 8725 John J. Kingman Road Fort Belvoir, VA 22060			10. SPONSOR/MONITOR'S ACRONYM(S) DTIC	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/ AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.				
13. SUPPLEMENTARY NOTES Autonomy - Human/Autonomous System Interaction and Collaboration; Scalable Teaming of Autonomous Systems; Machine Perception, Reasoning and Intelligence; Test, Evaluation, Validation, and Verification. C4ISR - HCI for Decision Making; Air Platforms – Unmanned Aircraft Systems (UAS)				
14. ABSTRACT <p>Artificial intelligence (AI) applied to weapons systems represents a major trend in research in the past 10 years. These initiatives seek to increase weapon accuracy, perform nonactive means of targeting, aid navigation and guidance and control (e.g., in Global Positioning System-denied situations), and reduce overall computational resources vs. traditional physics-based approaches to enable intelligent targeting on smaller, more affordable weapons systems. This research also includes extending the battlespace of operators to unmanned aerial vehicles and teaming with manned and unmanned platforms using swarming methods.</p> <p>We begin with an overview description and history of AI and outline the principals, techniques, and applications of AI for weapons systems. This includes a review of research and programs in supervising autonomous systems; guidance, navigation, and control; behavior and path planning; sensor and information fusion; intelligent strategy and planning; wargame modeling; and cognitive electronic warfare.</p> <p>We then offer a survey of systems and programs that apply AI for weapons systems. Although the focus is on U.S.-based systems and programs, a small subsection on related systems from Russia and China is included. We conclude with a brief commentary on the ethical considerations for using AI for weapons systems.</p>				
15. SUBJECT TERMS artificial intelligence, AI, weapons systems, unmanned aerial vehicles, swarm intelligence, deep learning, reinforcement learning, imitation learning				
16. SECURITY CLASSIFICATION OF: U		17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 92	19a. NAME OF RESPONSIBLE PERSON Vincent "Ted" Welsh
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED		19b. TELEPHONE NUMBER (include area code) 443-360-4600

THE AUTHOR

SAM CHAKOUR

Sam Chakour is a senior robotics software engineer (SW) specializing in autonomous systems and control, modeling and simulation (M&S) analysis, and model-based mission SW development. He currently works for QinetiQ Inc., in Lorton, VA, as a research and development engineer in the intersection of machine learning, computer vision, and signal processing for intelligence, surveillance, and reconnaissance systems and unmanned ground vehicle sensors. Prior to QinetiQ, he was a senior M&S software architect and M&S data analyst for autonomous unmanned aerial systems at Aurora Flight Sciences in Manassas, VA. He holds an M.S. in electrical engineering and an M.S. in applied mathematics from Grenoble-Alpes University and a B.S. in electrical engineering from George Mason University.

ABSTRACT

Artificial intelligence (AI) applied to weapons systems represents a major trend in research in the past 10 years. These initiatives seek to increase weapon accuracy, perform nonactive means of targeting, aid navigation and guidance and control (e.g., in Global Positioning System-denied situations), and reduce overall computational resources vs. traditional physics-based approaches to enable intelligent targeting on smaller, more affordable weapons systems. This research also includes extending the battlespace of operators to unmanned aerial vehicles and teaming with manned and unmanned platforms using swarming methods.

We begin with an overview description and history of AI and outline the principals, techniques, and applications of AI for weapons systems. This includes a review of research and programs in supervising autonomous systems; guidance, navigation, and control; behavior and path planning; sensor and information fusion; intelligent strategy and planning; wargame modeling; and cognitive electronic warfare.

We then offer a survey of systems and programs that apply AI for weapons systems. Although the focus is on U.S.-based systems and programs, a small subsection on related systems from Russia and China is included. We conclude with a brief commentary on the ethical considerations for using AI for weapons systems.



ACKNOWLEDGMENTS

We would like to extend a note of thanks to a few gracious subject matter experts (SMEs) who took the time to review a rough draft of this report and provide feedback. We appreciate their time and contributions.

Though some of our SMEs wish to remain anonymous, some agreed to be named. They are listed as follows:

- Laura Hiatt (Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory)
- Prof. Krishna R. Pattipati (Electrical and Computer Engineering Department, University of Connecticut, Storrs, CT)
- Nisar Ahmed (University of Colorado Boulder)
- Bradley Hayes (University of Colorado Boulder)
- Steve McGuire (University of California Santa Cruz)

CONTENTS

	ABOUT DSIAC	iv
	THE AUTHOR	vi
	ABSTRACT	vii
	ACKNOWLEDGMENTS	viii
	CONTENTS	ix
SECTION 1	INTRODUCTION	1-1
1.1	Problem Statement.....	1-1
1.2	Conventional Weapons Systems.....	1-1
1.3	Brief History of AI.....	1-2
1.4	What Is AI?.....	1-4
1.4.1	ANI.....	1-5
1.4.2	AGI.....	1-5
1.4.3	ASI.....	1-5
1.5	ML.....	1-6
1.5.1	Supervised Learning.....	1-6
1.5.2	Unsupervised Learning.....	1-6
1.5.3	Semisupervised Learning.....	1-6
1.5.4	Reinforcement Learning.....	1-6
SECTION 2	STATE-OF-THE-ART METHODS	2-1
2.1	Learning AI paradigm.....	2-1
2.1.1	Deep Learning.....	2-1
2.1.2	Reinforcement Learning.....	2-6
2.2	Stochastic Optimization and Search Algorithms.....	2-8
2.2.1	Stochastic Optimization.....	2-8
2.2.2	Graph Search Algorithms.....	2-11
2.3	Emerging AI Paradigms.....	2-12
2.3.1	Neurosymbolic AI.....	2-12
2.3.2	NE.....	2-12

CONTENTS, continued

SECTION 3	APPLICATIONS OF AI FOR WEAPONS SYSTEMS	3-1
3.1	Autonomy	3-1
3.1.1	Definition, Levels, and Frameworks	3-1
3.1.2	Autonomous System's Functional Components	3-4
3.2	AI in Perception	3-6
3.2.1	Image Segmentation	3-7
3.2.2	Target Detection, Classification, and Scene Understanding	3-7
3.2.3	Sensor Fusion	3-9
3.3	AI in Guidance, Navigation, and Control	3-9
3.3.1	GN&C Systems	3-10
3.3.2	Conventional Control-Theoretic Methods	3-10
3.3.3	Intelligent Control	3-11
3.3.4	Localization and Navigation	3-14
3.3.5	System Identification	3-15
3.4	Mission and Path Planning	3-16
3.4.1	GAs	3-17
3.4.2	Swarm Intelligence	3-17
3.5	Intelligent Strategy	3-19
3.6	Opponent Modeling and Wargames	3-21
3.7	Cognitive Electronic Warfare	3-24
3.7.1	Electronic Support Measures	3-24
3.7.2	ECMs	3-25
3.7.3	ECCMs	3-26
SECTION 4	SYSTEMS AND PROGRAMS APPLYING AI FOR WEAPONS SYSTEMS	4-1
4.1	Aerial Systems	4-2
4.1.1	Next-Generation Air Dominance Programs	4-2
4.1.2	Shield AI Hivemind	4-2
4.1.3	Shield AI V-Bat	4-2
4.1.4	Kratos XQ-58 Valkyrie	4-3
4.1.5	MQ-20 Avenger UCAS	4-3
4.1.6	Autonomous Loitering Munitions	4-3

CONTENTS, continued

4.1.7	Dynetics X-61 Gremlins.....	4-3
4.2	Maritime Systems.....	4-3
4.3	Land Systems.....	4-4
4.3.1	QinetiQ/Pratt Miller’s Expeditionary Autonomous Modular Vehicle.....	4-4
4.3.2	Textron Systems’ Ripsaw M5.....	4-4
4.3.3	Rheinmetall’s Lynx KF41.....	4-4
4.4	Swarm Systems.....	4-4
4.4.1	DARPA’s Offensive Swarm-Enabled Tactics.....	4-5
4.4.2	Collaborative Small-Diameter Bomb Swarm.....	4-5
4.4.3	Perdix Swarm.....	4-5
4.4.4	Mako UTAP22.....	4-5
4.4.5	Coyote UAS Block 3.....	4-6
4.4.6	Control Architecture for Robotic Agent Command and Sensing Swarm.....	4-6
4.4.7	Riptide Swarm Micro-UUVs.....	4-6
4.5	Battle Management and Intelligent Command & Control.....	4-6
4.6	ISR & Targeting Systems.....	4-7
4.6.1	SRC’s HPEC Pod.....	4-7
4.6.2	Nemesis.....	4-8
4.7	Navigation.....	4-9
SECTION 5	AI IN FUTURE COMBAT.....	5-1
SECTION 6	AI AND FOREIGN THREATS.....	6-1
6.1	Russia.....	6-1
6.2	China.....	6-2
SECTION 7	ETHICAL CONSIDERATIONS.....	7-1
SECTION 8	SUMMARY.....	8-1
	REFERENCES.....	9-1
	BIBLIOGRAPHY.....	10-1

CONTENTS, continued

FIGURES

Figure 1-1	AI History Timeline.....	1-2
Figure 1-2	AI Tiers.....	1-6
Figure 2-1	Nonlinear Model of Neuron in a Feedforward NN.....	2-2
Figure 2-2	Deep Feed Forward Networks.....	2-3
Figure 2-3	RL Architecture.....	2-7
Figure 2-4	DARPA's Air Combat Evolution Program Using AlphaDogfight in Simulations.....	2-8
Figure 2-5	Swarm Algorithm Concept for Each Agent.....	2-9
Figure 3-1	DoD 3000.09 Definitions of Top Levels of Autonomy.....	3-3
Figure 3-2	Autonomy Functional Component.....	3-5
Figure 3-3	Semantic vs. Instance vs. Panoptic Segmentations.....	3-7
Figure 3-4	MLP-Based NARX Architecture.....	3-16
Figure 3-5	Air Combat Evolution.....	3-16
Figure 3-6	EW System Components.....	3-24
Figure 4-1	V-Bat UAV in Transition Flight From Vertical Takeoff.....	4-3
Figure 4-2	U.S. Air Force's Skyborg Conceptual Design.....	4-3
Figure 4-3	Dynetics X-61 Gremlins UCAS.....	4-4
Figure 4-4	U.S. Navy's Sea Hunter in Its First Demonstration.....	4-4
Figure 4-5	QinetiQ/Pratt Miller's RCV-L.....	4-4
Figure 4-6	Army's Ripsaw M5 Unmanned Battle Vehicle.....	4-5
Figure 4-7	Swarm II USVs With CARACaS During Demonstrations.....	4-6
Figure 4-8	Riptide Autonomous Solutions Micro-UUVs.....	4-7
Figure 4-9	MQ-9 Reaper in Flight.....	4-8
Figure 4-10	HRL's Nemesis Eye Interprets the Battlefield for Actionable Intelligence.....	4-8
Figure 5-1	The Army's "Robotic and Autonomous Systems Strategy" Foresees Soldier/ Robot Teaming and Robot/Robot Teaming as Strategic Mission Enablers.....	5-1
Figure 7-1	The DIU's RAI Guidelines Workflow During the Planning Phase of the AI Life Cycle.....	7-3
Figure 7-2	The DIU's RAI Guidelines Workflow During the Development Phase of the AI Life Cycle.....	7-3
Figure 7-3	The DIU's RAI Guidelines Workflow During the Deployment Phase of the AI Life Cycle.....	7-3

TABLES

Table 3-1	SP&A-Based Levels of Autonomy.....	3-4
-----------	------------------------------------	-----

SECTION 01

INTRODUCTION

1.1 PROBLEM STATEMENT

The recent advances in machine learning (ML) and AI research shed light on the power and potential of AI in enabling innovations, increasing the utility of machines, and augmenting the human capability and experience. The disruptive nature of AI technologies and the depth of their impact are yet to be fully grasped by the general public at large. Considering the new era of emerging technological threats, it is important to showcase the critical and relevant AI research and state-of-the-art technologies that not only provide weapons systems with increasing autonomy over conventional weapons systems but significantly increase their lethality and combat survivability. Ultimately, AI presents colossal and strategic opportunities in developing game-changing technologies that will ensure our national security, prosperity, and technological leadership.

1.2 CONVENTIONAL WEAPONS SYSTEMS

The U.S. military has made monumental strides in creating advanced, conventional weapons technologies that support the missions and enhance the capabilities of our soldiers on the battlefield. These conventional weapons technologies are mostly automated systems that rely on a preprogrammed set of rules in planning, executing, and accomplishing a task or mission. However, on the frontier end of the newly developed weapons of countries like China and Russia, AI-enabled warfare and hypersonic weapons pose a new breed of qualitative

challenges for the U.S. Armed Forces. The pace of the next-generation combat requiring time-critical and copious combat information processing for strategic decision making consigns many of the U.S. conventional weapons systems to low-risk missions and a posture of diminished deterrence outside the nuclear realm.

It must be acknowledged that humans alone are expensive assets to train. Adding more personnel to the battlefield is not an elegant or cheap solution to advancing the state-of-the-art of war. Instead, augmenting human-in-the-loop systems with AI-enabled intelligent hardware can provide more eyes and ears on the battlespace and free up human decision making by enabling AI systems to perform some tasks that are easy and routine.

Furthermore, unmanned combat aerial systems (UCASs) are a proven cost-effective systems solution for intelligence, surveillance, and reconnaissance (ISR) missions and remote airstrikes. However, the automated capabilities are still bounded by human-in-the-loop operations, evaluation, and engagement. While there is no intent on eliminating the human element in weaponized AI systems in any foreseeable future, the capability of humans continues to constitute a ceiling in the synergetic potential of these systems. But a new ecosystem of AI-enabled intelligent weapons systems would usher in new forms and strategies of warfare.

In its 2021 report, the National Security Commission on Artificial Intelligence submitted that the U.S.

Department of Defense (DoD) military enterprise trails behind the commercial sector in terms of integrating AI-enabled technologies and urged that the foundations for widespread integration of AI across the DoD be in place by 2025 [1].

1.3 BRIEF HISTORY OF AI

The concept of artificially replicating a facet of human intelligence in some form was contemplated by philosophers for centuries. In 1869, William Jenon created the first machine that implemented logic computation based on Boolean logic. The machine was capable of computing Boolean algebra and Venn diagrams faster than humans. With this development of logic-computing machines, it was natural to question whether machines could reason through logic to solve problems and make decisions for humans. An illustration of the history and evolution of AI is shown in the timeline in Figure 1-1 and expanded on in this section [2].

In some of the earliest work in theoretical computer science, British mathematician Alan Turing pondered the question of whether machines can behave and solve problems intelligently and in the same manner as humans. He posited in his Turing test that if a machine mimics indistinctly an intelligent being such as a human, then the machine is intelligent. This theoretical test became a guiding formalism in which current machines are tested for their capacity or potential to mimic the human concept of intelligence. As a testimony to the test, the Loebner prize is a Turing test competition whose mission is to evaluate the current state of machine intelligence research against the fundamental question posed by Turing.

In 1928, John von Neumann proved the theorem underlying the Minimax algorithm, which seeks to provide a strategy for minimizing the maximum possible loss during zero-sum games play.

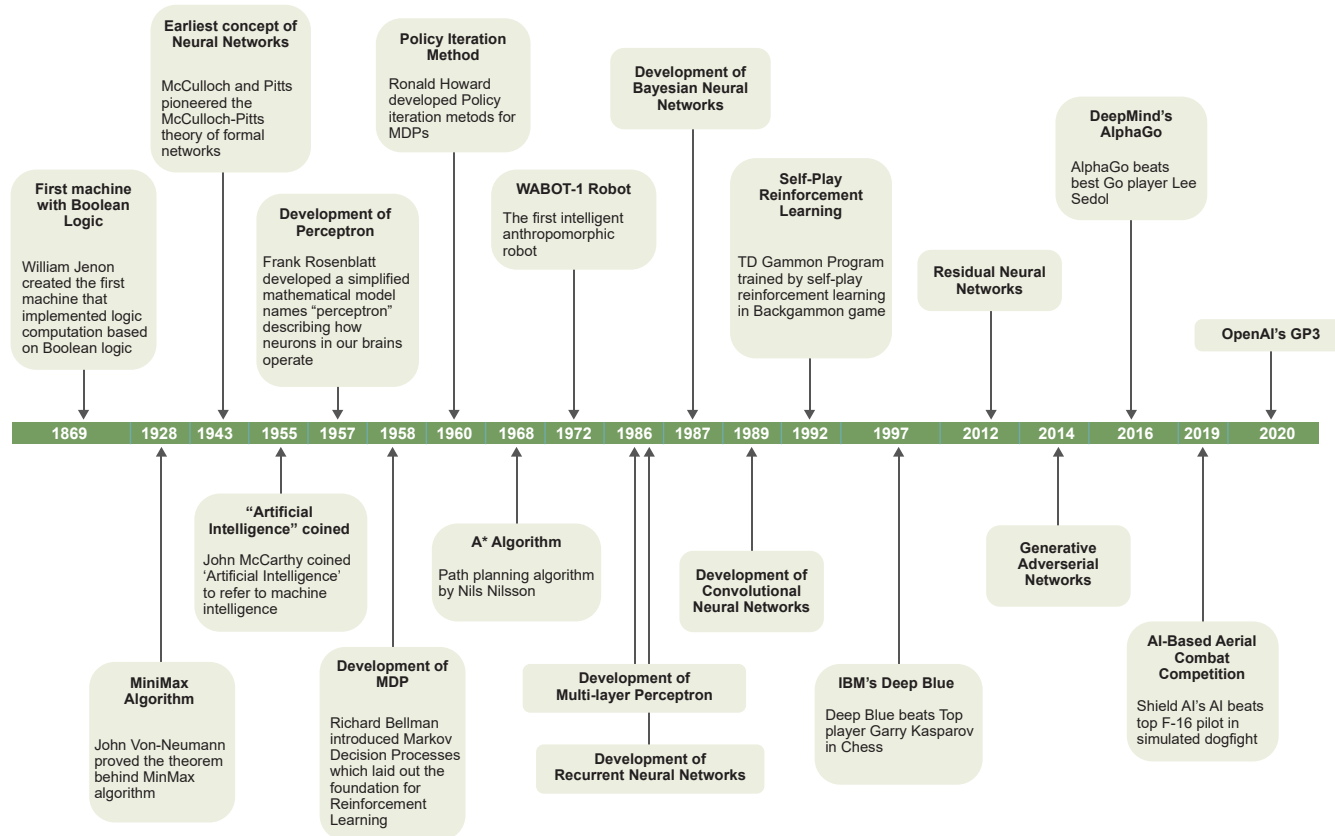


Figure 1-1. AI History Timeline (Source: QinetiQ).

During the height of the Second World War, Alan Turing and his teams developed a machine algorithm that could decipher the German Enigma message codes. The success of his algorithm, which fueled further endeavors in delegating complex tasks to a machine, was the foundation to machine computation and a precursor to the development of ML.

In 1943, McCulloch and Pitts pioneered the earliest concept of neural networks (NNs)—McCulloch-Pitts Theory of Formal Networks—which was featured in four of von Neuman’s lectures at the University of Illinois in 1949 [3].

Around the same time, John McCarthy, a computer scientist who coined “artificial intelligence” to refer to machine intelligence in 1955; computer scientist Allen Newell; and Herbert A. Simon, an economist and political scientist, pioneered the first true program aimed at automated reasoning (called the Logic Theorist). With this ground-breaking effort, the quest for intelligent machines began, paving the way for AI as a new field of academic research in computer science.

In 1957, a psychologist named Dr. Frank Rosenblatt developed a simplified mathematical model named “Perceptron,” which described how neurons in our brains operate. The achievement was highlighted as the “Perceptron Convergence Theorem.”

That same year, Richard Bellman developed dynamic programming for solving a class of optimal control problems. He also introduced the Markov Decision Processes formulation of discrete stochastic optimal control problems, which laid out an important foundation for what is now referred to as “reinforcement learning.”

Following these developments, another AI pioneer named Arthur Samuel successfully developed the first checkers algorithm using his earlier seminal work in ML. He had implemented an early version of what is now known as “Alpha-Beta Pruning,” which is a search tree method that reduces the

number of evaluated nodes by the Minimax algorithm. An early version of a nonparametric, supervised learning method called Decision Trees was developed by a statistician named William Belson in 1959.

In the 1960s, AI research focused on solving mathematical and optimization problems. The policy iteration method for Markov Decision Processes was proposed by Ronald Howard in 1960, establishing some of the earliest work related to reinforcement learning.

By 1968, the well-known path search algorithm called A-star was proposed by the computer scientist Nils Nilsson. Advances in robotics modeling, control, and machine vision were made in the late 60s, leading to the development of the first “intelligent” anthropomorphic robot named WABOT-1 in 1972 and integrating limb manipulation, vision, and speech systems.

The revival of the Harry Klopf’s “Heterostatic Theory of Adaptive Systems” was influential in the development of trial-and-error paradigm of adaptive systems. In 1977, Ian Witten proposed one of the earliest reinforcement learning systems that used temporal-difference methods. Richard Sutton and Andrew Barto devised a reinforcement learning algorithm called the Actor-Critic Method.

Because of the computational limitations of the mid-70s to late 80s computers, AI research found difficulties in applications with large data processing requirements, such as in vision learning or optimization problems. At the same time, mathematical research “proved” that a (single layer) perceptron could not learn certain patterns. Furthermore, a Lighthill report published in 1973 was very pessimistic about the potential for AI, which caused funding to be cut for AI research. As a result, the funding shortage led to research in AI to experience a period known as “AI Winters.”

By the mid to late 80s, important theoretical contributions were made in NNs following the

development of the multilayer perceptrons in 1986. These contributions were the development of recurrent neural networks (RNNs) in 1986 by David Rumelhart, Bayesian Networks in 1987 by John Denker et al., and the convolutional neural networks (CNNs) by Yann LeCun in 1989. Moreover, Chris Watkins developed another important reinforcement learning method dubbed “Q-Learning” in 1989.

In 1992, at IBM’s Thomas J. Watson Research Center, Gerald Tesauro trained the TD Gammon program for the backgammon game by self-play reinforcement learning. IBM’s Deep Blue computer beat chess world champion Garry Kasparov in 1997 using brute-force, search-based algorithms, making it the first program to win against a top professional in chess.

During the late 90s and early 2000s, much of the progress seen in ML was driven by the exponential progress made in computer processing, storage, and distributed computing. In 2007, the guaranteed optimal play requiring significant computing resources was solved for checkers.

The surge of graphical processing unit use for general computing in the last two decades led to further progress in AI applications today, specifically, the development of the different NN topologies such as residual networks and generative adversarial networks in 2012 and 2014.

In 2015, the ImageNet competition, an open competition to develop a classifier for the ImageNet image set of some four million images, had a winner with an error rate considered to be lower than a single human. In 2016, DeepMind’s AlphaGo program became the best AlphaGo player after beating Lee Sodol, who was considered the best player of the Go game at the time. Following the learning capability of AlphaGo, AlphaZero extended AlphaGo in 2017 to become the best player at chess and Shogi.

In 2019, the U.S. Defense Advanced Research Projects Agency (DARPA) launched AlphaDogfight, a series of three competition rounds of AI-based aerial combat algorithms against top-trained pilots in a simulated F-16 dogfight. The first and second rounds of the competition entailed the AI programs competing against each other. The third round distilled the AI victor pilot to compete against the top U.S. Air Force Weapon School graduate. Heron System’s AI pilot not only won against the competing AI aerial combatants but secured an incredible score of five victories against the highly trained human F-16 pilot.

OpenAI introduced a “natural language processing” model called GPT-3 in May 2020, which generated writing content indistinguishable from a human. Its latest version can generate programming language code from simple descriptive language [4]. The history of AI continues forward, especially for DoD applications to weapons systems. The remainder of this report will survey contemporary AI techniques and systems related to weapons systems.

1.4 WHAT IS AI?

According to Barr and Feigenbaum, AI is defined as “the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit characteristics we associate with intelligence in human behavior – understanding language, learning, reasoning, solving problems, and so on” [5].

A more recent definition of AI is provided by Stuart Russel and Peter Norvig in their book *Artificial Intelligence: A Modern Approach* as “the designing and building of intelligent agents that receive percepts from the environment and take actions that affect that environment” [6].

Pei Wang elegantly defined intelligence as “adaptation with insufficient knowledge and resources” [7]. Although the definition does not state the purpose of adaptation (e.g., objective), it sheds light on what is needed to be accomplished to reach this intelligence.

If AI is to be defined anthropocentrically, i.e., to perform tasks at the human-level of intelligence, then an AI requires perception, reasoning, knowledge construction, inference, decision and planning, learning, communication, and the ability to move and manipulate the environment efficiently.

AI's scientific goal is to answer which ideas about knowledge representation, learning, rule systems, search, etc., explain various sorts and levels of real intelligence. The engineering goal is to develop AI techniques for the different domains of application to tackle real-world problems.

At the root of AI's scientific foundation, we find identifiable concepts from different scientific fields—philosophy, logic/mathematics, computation, psychology and cognitive science, biology and neuroscience, and evolution. Contributions from these different domains of knowledge already proved inevitable and indispensable in the quest for discovery and better understanding of what AI is or would be. Many fields researching AI are concurrently constructing models of how human cognition operates and adopting useful concepts between them. For example, the NN, a concept which originated from biology, attempted to build artificial systems based on simplified artificial neurons, a concept which led to the representation of a simple abstract knowledge structure powerful enough to solve large sets of computational problems.

AI is broadly classified into three main tiers—Artificial Narrow Intelligence (ANI), Artificial General Intelligence (AGI), and Artificial Super Intelligence (ASI). Figure 1-2 illustrates the various groupings within the three tiers, which are discussed more in this section.

1.4.1 ANI

ANI is a description of an AI system that performs a narrow or singular task. It can include various methods to obtain the result, such as traditional

ML (image classification as an example) or target detection (both ML and rule-based systems). Given a set of rules or constraints, its objective is to provide a set of outputs representing a narrow set of tasks. ANI does not expand or learn from new perception nor does it self-learn a new mode of operation. Data mining, most expert systems, and predictive functions specific to an application (e.g., spam detection and facial recognition) are all considered forms of ANI. ANI would also include “limited memory AI”—the type of system used in self-driving cars, using past experiences (training), and learning to make decisions and improve over time.

1.4.2 AGI

AGI is a stronger form of intelligence, as it is augmented by more human intelligence-like traits, such as the ability to learn on its own and interpret emotions and speech tone. This places the intelligence associated with AGI on par with a human's level of intelligence. Some key core capabilities of AGI are as follows:

- The ability to reason, solve problems, employ strategy, and make decisions under uncertainty.
- The ability to show knowledge.
- The ability to plan.
- The capability to learn.
- The ability to communicate in a natural language.
- The ability to integrate all the above toward a common goal.
- The combination of human-like thinking with computation like the Turing test.

1.4.3 ASI

ASI models an intelligence that surpasses the brightest human minds. Methods for achieving ASI are still being conceptualized but would be those systems that go beyond AGI and entail some sort

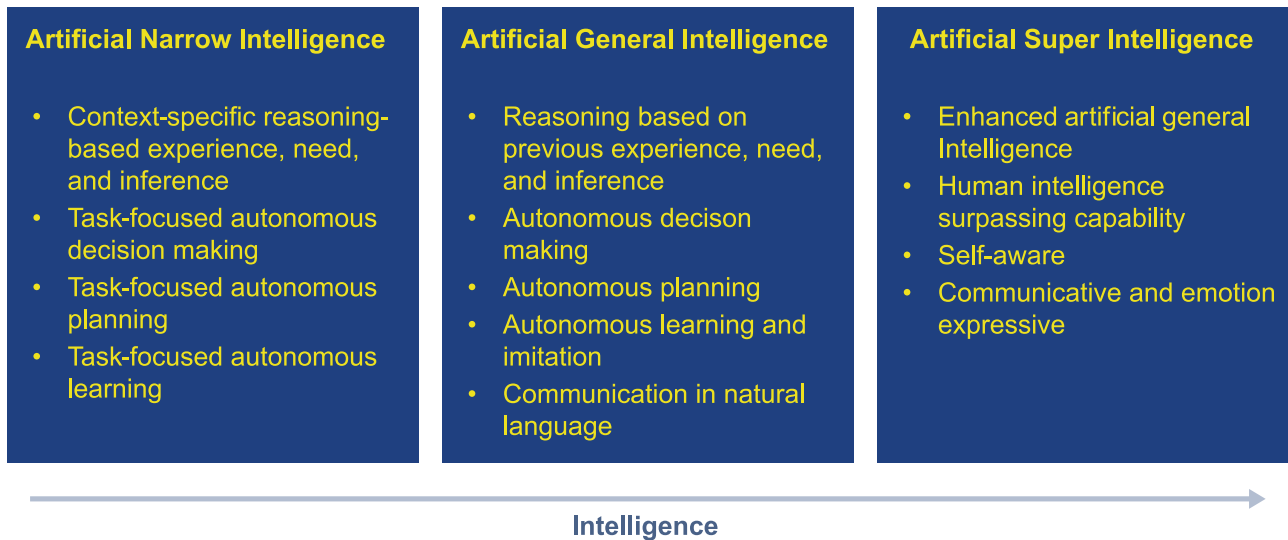


Figure 1-2. AI Tiers (Source: QinetiQ).

of self-awareness. These systems would ideally represent all human cognitive capability and more.

1.5 ML

ML is the ability of a machine to learn from data for the purpose of making accurate predictions. It is broadly divided into four classes of learning that provide a rich family of dedicated and generalized techniques.

1.5.1 Supervised Learning

In this form of learning, training data use contained input and labeled or predefined output data. If there are missing input or output entries, they are preprocessed so that an input is mapped properly to its true corresponding output. By learning from the properly generated training dataset, the system learns to associate an input not in the original dataset to its predicted output (label or value). Typical problems addressed by this type of training are regression and classification [8].

1.5.2 Unsupervised Learning

The system in this form of learning discovers interesting or hidden structures directly from the unlabeled data [9]. The unsupervised learning is

used for cluster analysis, dimensionality reduction, or estimating the density likely to have generated the input data [8].

1.5.3 Semisupervised Learning

When the dataset contains labeled and unlabeled data, the system in this form of learning makes use of the unlabeled data to better capture the underlying data distribution and obtain a better prediction had it trained from the labeled data alone. This form of learning is suitable in situations when there are far less labeled data than the unlabeled one in the training dataset [8].

1.5.4 Reinforcement Learning

In this mode of learning, the system trains using a reward/penalty mechanism such that it selects and performs actions that either lead the system to receive rewards when the actions are desirable or receive penalties when the actions are not desirable. Reinforcement learning problems involve learning what to do (how to map situations to actions) to maximize a numerical reward signal [9].

SECTION 02

STATE-OF-THE-ART METHODS

This report focuses on specific AI methods central to developing state-of-the-art technologies relevant to weapons systems, particularly in ML-based AI (or simply learning AI) and advanced stochastic optimization methods in solving problems related to the behavior of autonomous systems in a complex, dynamic environment. The learning-AI paradigm has eclipsed the earlier paradigm of knowledge-based AI systems; hence, the latter is not explored in this report. The techniques presented do not represent an exhaustive survey within their respective paradigms but instead represent either fundamental methods proven in state-of-the-art robotics and autonomous systems, an amelioration of fundamental techniques with interesting properties, or recent techniques having potential application in weapons technologies.

2.1 LEARNING AI PARADIGM

2.1.1 Deep Learning

Deep learning refers to a subset of powerful algorithms within ML that employs deeply connected artificial neural networks (ANNs) called deep neural networks (DNNs).

According to Yoshua Bengio [10]:

“Deep learning algorithms seek to exploit the unknown structure in the input distribution in order to discover good representations, often at multiple levels, with higher-level learned features defined in terms of lower-level features.”

The building block of a DNN is the perceptron, which consists of an information processing unit called an “artificial” neuron. Inspired by the structure and behavior of the neurons in the brain, a neuron provides the fundamental operation of an NN.

The modern concept of an “artificial” neuron has the following three basic elements [3]:

1. Synapses: These are connecting links weighted or characterized by a value strength.
2. Transfer Function: This function traditionally represents a linear regression model (linear combiner) summing the input signals weighted by their respective weights or strength values.
3. Activation Function: A function for limiting the amplitude of the output of the neuron. The transfer function may include an externally applied bias value that modulates (increases/decreases) the net input of the activation function.

The perceptron is the simplest form of an NN formed by a single neuron with adjustable synaptic weights and a bias [3].

The multilayer perceptron (MLP), as a generalization of the single-layer perceptron and fundamental feedforward network architecture (Figure 2-1), owes its success to the error backpropagation algorithm, which is based on the error-correction learning rule. Using a gradient-descent-based method and the chain rule, the algorithm aims to minimize a loss function (e.g.,

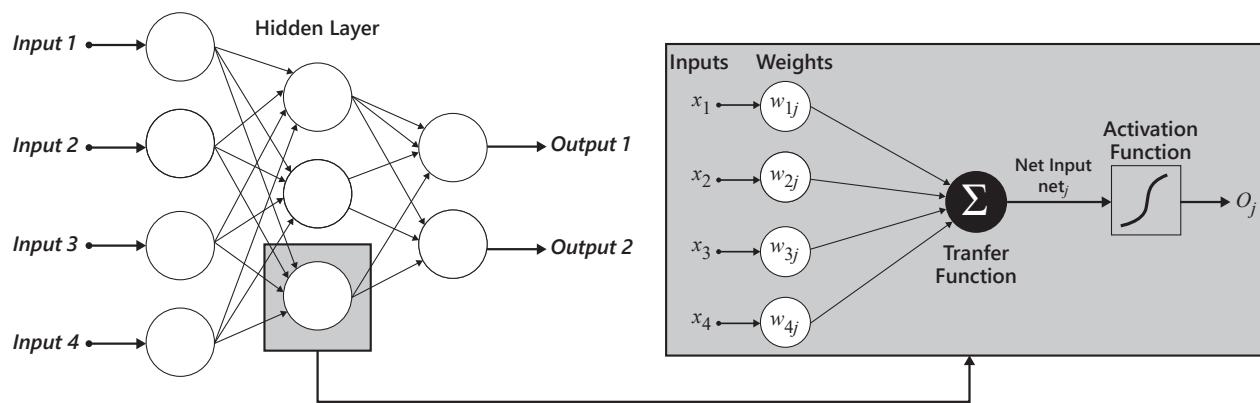


Figure 2-1. Nonlinear Model of Neuron in a Feedforward NN (Single Hidden Layer) (Source NCSAI [1]).

squared error, cross entropy, etc.) by propagating its change regarding the free parameters (weights and biases) backward through the network against the direction of the synaptic connections [3].

A DNN is deep due to the hidden layers not visible in the input or output nodes. The depth refers to the layers in the network, whether it has a feedforward, feedback, or convolutional architecture. There is no general rule to the number of hidden layers an NN needs to have to be a DNN. It is, however, accepted that an NN with three or more layers (including the output) having nonlinear activation function is a DNN (Figure 2-2).

While DNNs have better performance overall than shallow NNs, their depth is a source of several issues, such as the following:

- **Vanishing Gradient:** This problem is encountered when computing the gradients in the backpropagation algorithm during the learning process of an NN. During the partial derivative computation of the loss function regarding each free parameter, it may occur that when the gradient of the bounded activation function (e.g., hyperbolic tangent is bounded in $[-1, 1]$ and logistic function is bounded in $[0, 1]$) becomes small, the chain rule of the gradient values leads to a vanishingly small number in networks with a large number of layers, preventing the improvement of the learning process [11]. The problem is often

addressed by nonbounded activation functions (e.g., rectified linear activation function [ReLU]) or renormalization.

- **Degradation:** This problem refers to the increase in error rate as the number of hidden layers of an NN is increased, thus affecting both the learning ability and the informative capacity of an NN.
- **Overfitting:** This problem occurs when a network model becomes so overly fitted to the training dataset that it no longer applies to the novel dataset. The model learns the details, noise, and outliers in the training dataset, thus negatively affecting its performance. The problem may occur also in complex network models with many hidden layers.
- **Hard to Train:** As the network structure becomes more complex, the gradient-based backpropagation algorithm, which is an unconstrained optimization problem, becomes difficult to solve due to nonconvexity of the loss-function hyperplane and the NP-complete nature of the optimization problem [12].

DNNs are supervised learning techniques, but some network topologies like autoencoders operate in an unsupervised manner. They can also be used to cluster input based on similarities. One can extract the features with an NN and then deploy an unsupervised methodology such as k-means clustering, which makes the NN architecture a semisupervised DNN.




Figure 2-2. Deep Feed Forward Networks (Source: NCSAI [1]).

The following fundamental and important topologies represent different deep-learning model structures:

1. MLP: The common class of feedforward ANN consisting of at least three layers—input, hidden, and output—and nonlinearly activating nodes.
2. RNN: A class of NNs that maintains the temporal context of data in its internal memory to predict time-dependent dynamics in sequential data. They use feedback connection to hold internal states as memory so sequence processes can be remembered, which makes them different than the feedforward NNs [13].
3. Long Short-Term Memory (LSTM) Networks: An RNN variation devised to overcome the computational limitation called “vanishing gradient” of RNNs due to the long-term dependency problem of the latter. The LSTM network avoids the long-term dependency problem by retaining the longer-term context with the previous prediction [14]. LSTM networks are used for time-series processing, speech recognition and synthesis, audio processing, and financial forecasting.
4. Autoencoders: A class of NNs used for the unsupervised learning of data representation via encoding, compression, or reconstruction. Autoencoders are considered self-supervised, learning NNs and are used in dimensionality reduction, features clustering, and data compression [15].
5. Nonlinear Autoregressive eXogeneous (NARX) Network: A recurrent dynamic network with feedback connections enclosing several layers of the network. The network adopts the linear autoregressive model, which is used to model time-series dynamics. NARX can be used to predict the output of nonlinear dynamic systems (system identification) and for nonlinear filtering [16].
6. CNNs: An MLP designed specifically to recognize two-dimensional (2-D) shapes with a high degree of invariance to translation, scaling, skewing, and other distortions [3]. The core concept introduces hidden convolution and pooling layers to identify spatially localized features via a set of receptive fields in kernel form. These structures were devised to have important properties that are neurobiologically inspired, such as feature mapping and subsampling. As a result, CNNs extract multiscale, localized spatial features to construct high-level representations.
7. Gated Recurrent Unit (GRU): A recent alternative approach to solving the “vanishing gradient” problem of RNN; hence, it attempts to solve a similar class of problems pertaining to time-series processing. While LSTM is more accurate on a larger data set, a GRU uses less memory and runs faster due to its structure that contains less parameters. Therefore, it uses less training parameters overall to compute [17], which is a tradeoff justified in many applications.

8. Residual Neural Networks (ResNets): A recent development in deep CNN that introduces a scheme devised to solve several observed issues in deep CNNs (namely, the vanishing gradients and degradation problems [18]) and improve the performance of previous deep CNN architectures like Visual Geometry Group (VGG) architectures.

One main advantage of ResNets is that they are faster to train because inputs can forward propagate faster through the residual connections (or skip the connection, which provides another path for data to reach the latter part of the network by skipping layers) across the layers.

ResNets are used in object recognition, activity recognition, and scene understanding and have been also successfully used for NLP problems, such as machine comprehension (question/answer) and speech emotion recognition [19].

9. Densely Connected Networks (DenseNets): A type of CNN that leverages dense connection structure between layers through modules called Dense Blocks, where all the layers are connected directly with each other. Feedforward nature is maintained by propagating feature maps forward through the layers. DenseNet is cited to provide the best representation of images when compared to other CNN-based architectures [20] when applied to near-identical images in ImageNet (a large visual database for visual object recognition research). It is suggested that DenseNets are also easier to train compared to architectures of similar size [20, 21].

10. Sparse Network (SparseNet): A modified architecture to further improve the performance of DenseNet by sparsifying the density. This architecture offers comparable results to DenseNet while being smaller and faster (2.6× smaller and 3.7× faster) than the original DenseNet architecture [22]. The architecture does, however, require a layer depth between 28 and 76.

11. Generative Adversarial Networks (GANs): Networks used for unsupervised learning tasks. They are composed of one generative model, while the other is a discriminative model. The two models compete to generate samples from the statistical distribution of the original samples. As a result, the samples generated are closely matched to the original samples [23].

GANs are used to interpret content from data and create similar but novel versions of the data. However, they can suffer from convergence issues and, hence, are hard to train.

12. Graph Neural Networks (GNNs): An architecture based on graph data structure, which models a set of objects (nodes) and their relationships (edges). Motivated partly by the lack of CNNs to handle non-Euclidean types of data (e.g., unordered, nonuniform grid graphs), GNNs were created to perform inference on graph-represented data having complex relationships and object interdependencies [24, 25].

GNNs have been applied to numerous applications of structured and unstructured data, such as the following:

- Computer Vision: Image classification and scene graph generation.
- NLP: Semantics exploitation in machine translation, relation extraction, and question answering.
- Recommender Systems.
- Program Reasoning.

13. Deep Belief Networks (DBNs): A generative model NN devised to address the performance shortcoming of traditional deep-layered NNs—namely, the learning time, large number of data required for training, and the convergence issues due to inadequate selection of parameters. DBNs are based on generative, stochastic NNs that can learn a probability distribution from its inputs. They are constructed by restricted Boltzmann machines

or autoencoders, with the hidden layers of each subnetwork serving as the invisible layer for the next layer. They are trained in a greedy manner, are less computationally expensive compared to forward NNs, and suffer far less from the vanishing gradient problem [26].

The following drawbacks associated with DBN are numerous:

- Requires significant computation resources for a large dataset.
- Requires a solid theoretical grasp of their structure and how they function by the user.
- Requires classifiers.

DBNs are known to capture deep hierarchical representations of input features. They have been used for image recognition [27], motion-capture application [28], and nonlinear dimensionality reductions [29].

Because DNNs, in general, are computationally harder to train, classic optimization methods for the backpropagation “Error Propagation” algorithm can have convergence issues. State-of-the-art optimization methods for the backpropagation algorithm in DNNs include the Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm and conjugate gradient algorithms. These algorithms are formulated to adapt their computation by tuning the learning rate and step-size parameters and compute derivative-based information (e.g., Hessian matrix) to home-in on the valleys in the search hyperspace for optimal solutions of the free parameters.

The deep-learning ability to generate complicated models from large data leads to improved accuracy and expressiveness. As a result, deep learning provides the following set of pertinent solutions to a wide range of application within AI problem domains:

- Recognition in speech, images, and video.
- Context evaluation in multidimensional data

(e.g., imaging, videos, radar, and monitoring systems).

- Time series forecasting.
- Function approximation.
- Data compression.
- Data outliers’ detection.
- System identification.
- Scene understanding.
- Explainable AI.
- Autonomous driving.

Deep learning provides the following advantages over classic ML algorithms:

- Scalability supported by parallel and distributed computing.
- Complex feature extraction.
- Hierarchical feature learning.
- Ability to deal with unstructured data.
- Adaptable DNN architecture to various types of problems.

Some known limitations of deep learning are as follows:

- Incorporating logic in hybrid topologies is difficult to achieve.
- Training time may be extremely long for complex models or high-dimensional feature space.
- Generally, massive training datasets for convergence are required.
- Convergence is difficult for some structures.
- No theoretical formalism exists to guide the selection of the proper topology and parameter tuning for a given problem. The selection is often a domain knowledge.

Deep learning provides a powerful framework of techniques for hierarchical learning to find multiple high-level abstract representations of patterns in complex data.

2.1.2 Reinforcement Learning

The concept of reinforcement originated in the work of psychologist Ivan Pavlov in 1903 when he observed the conditioning process of animals via a reinforcement mechanism [30]. The mathematical foundation of the problem, however, was introduced by Richard Bellman in 1957 in his work when he formulated the problem of deciding best actions that meet a cost function under uncertainty as the Markov Decision Process (MDP).

2.1.2.1 MDP

MDP is a discrete-time stochastic process for modeling sequential decision making of an action-based agent in a stochastic environment in which this agent can act and where the outcome of the action is uncertain [30].

2.1.2.2 Partial Observation MDP

The Partial Observation MDP (POMDP) models the case in which uncertainty is extended to the state. Instead of exact observation of the state, only the probabilistic relationship with the state is observed. A common solution to POMDPs is the inference of a belief distribution over the underlying state at the current time step and then applying a policy that maps beliefs into actions [30].

Reinforcement learning (RL) models the natural learning process of humans, animals, and several biological systems via the action-reward mechanism. RL is categorized in two forms—single-agent RL (SARL) and multiagent RL (MARL).

2.1.2.3 SARL

In SARL, the RL model is divided into states, actions, and rewards (Figure 2-3). States are a representation of the current world or environment of the task. Actions are something an RL agent can do to change these states, while rewards are the utility the agent receives for performing the desired actions. The states inform the agent the direction

toward a goal via the rewards signal. The objective is to learn a policy that specifies through time which actions to take from each state to maximize the cumulative reward [9].

The reward signal received by the agent determines the quality of the action. Unlike in supervised learning (SL), where the actions are directly mapped to desired behavior, the feedback “reward” in RL is less informative [31]. The agent must balance its exploration of its environment by relying on its previous knowledge of actions and rewards to take uncertain future actions with unknown rewards or penalties in the hopes it discovers more rewarding actions. This balance of discovery and action is referred to as exploration-exploitation tradeoff [32].

Unlike the control theoretic approach, the environment in RL is represented by all aspects external to the agent control—systems dynamics, process errors, nonlinearities (e.g., saturation, dead zones, and hysteresis), delay, sensors’ noise, disturbances (e.g., steady wind, gusts, center of gravity, and the moment of inertia displacement), and any variations from the nominal values. The environment also is represented by the uncertainties related to high-level decision logic. Furthermore, the logical structure is embedded in the control authority represented by the policy function (controller).

One of the key advantages to RL is that it can be model free (system dynamics representation) and does not require knowledge of the reward function [33]. The formulation of the system’s state of evolution in the environment as an MDP allows generality of the type of problems RL can solve, even a nonquadratic reward function (cost function) and a system with stochastic nonlinear dynamics. The control structures that result from the learning process are generalized structures capable of adapting to changes in systems dynamics and the environment (the nature and severity of disturbances in the control-theoretic context) while automatically acquiring suitable

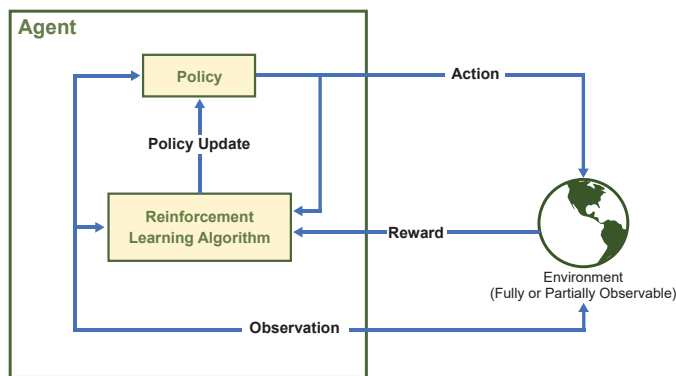


Figure 2-3. RL Architecture (Source: QinetiQ).

control policies that improve through the learning process. This makes RL less restrictive, extensible to other similar systems (with different dynamics) without change to the control scheme, and scalable [34], implying its applicability to a large set of complex problems requires strategy and adaptation.

2.1.2.4 MARL

MARL is another form of RL that differs from SARL. A robotic manipulator learning to organize its environment, a humanoid learning to walk, or a car learning to park on its own are all examples of SARL. MARL is concerned with how multiple agents interact with one another and their environment [35, 36]. MARL has two theoretical frameworks—stochastic games and extensive-form games.

Stochastic games can be construed as a multiagent generalization of MDP. Depending on the objective, MARL may address cooperation, competition, adversariness, or a combination of these settings. Robotic rescue and collaborative manufacturing are examples of cooperative MARL, while swarm unmanned aerial systems (UASs) are examples of mixed MARL. Under the stochastic games, MARL is accepted to be limited to a fully observed state space in which the agent has perfect information on the system state and executed actions at given time step [35]. In a competitive setting, namely zero-sum and constant-sum utility, MARL has been demonstrated

in an imperfect information setting for multiagent decision making.

The distributed nature of MARL comes with several advantages—experience sharing between agents, imitated behavior of a human or qualified agent, and redundancy inherent in homogeneous MARL, which provides increased level of robustness. However, the multidimensionality of the state space and the combinatorial nature of MARL pose several challenges as follows [30, 31]:

- **Nonstationarity Problem:** The environment may become nonstationary as each agent acts on it to improve its policy concurrently.
- **Nonuniqueness of the Learning Goal:** The multiobjective criteria of each agent may not be aligned with those of other agents and, therefore, consensus equilibrium points are not reached.
- **Scalability Problem:** As the number of agents increases in any MARL setting, the joint action space increases exponentially. This is known as the curse of dimensionality.

Also note that scalability and imperfect information are open problems in both SARL and MARL due to the large state and action spaces and their continuity.

2.1.2.5 Deep Reinforcement Learning (DRL)

DRL is the result of the recent progress in RL research that demonstrates the benefits of incorporating DNNs in RL algorithms. DNNs serve as function approximators of states to values or state-action pairs to an action-value criterion. When used with RL, DNNs interpret inputs and provide prediction of policy [37].

State-of-the-art, model-free RL algorithms are Deep Q-learning (Dueling Deep Q-Network [DQN], Prioritized Dueling DQN), Policy Gradient Method (Actor-Critic, Advantage Actor-Critic, Asynchronous Actor-Critic, and Actor-Critic With Experience Replay), Deep Deterministic Policy Gradient, Soft

Actor-Critic (a bridge between Q-learning and DPG). State-of-the-art model-based algorithms are Manifest State Based and Latent State Based [38].

2.1.2.6 Demonstrations of RL

In the last decade, pivotal state-of-the-art AI developments in RL by research- and industry-leading technology firms have captured a wide interest from ML practitioners across different fields of application and sectors. Examples of RL's success are described next.

DeepMind's AI bots reached super-human performance after defeating the top human professionals in Go, chess, and Shogi games. OpenAI developed an AI that defeated top players in a real-time strategy, multiplayer game environment and developed an RL-based motion control algorithm for an anthropomorphic robotic hand with perception and motion dexterity to solve Rubik's Cube [39]. DARPA's AlphaDogfight competition (Figure 2-4) ended with a crushing

win for Heron Systems AI against the human pilot, ushering an era of credible, autonomous, and potentially lethal machines. DRL success was further demonstrated in other atypical applications, such as the automated reconfiguration of network topology of software-defined networks to address network traffic fluctuations and minimize network delay [40].

When using RL for defense applications, robots may operate with a human in the loop to accept or reject the recommended actions of the RL algorithm. Much like early self-driving cars operate with a backup human driver to take over when the algorithm's actions were hazardous, a robot with a human in the loop may be necessary and practical in situations where human life is actually or potentially at stake. In this case, the robot should be thought of as an assistive device rather than a fully autonomous vehicle. The human operator then retains the capacity and responsibility to veto or override the robot's decisions as needed.

2.2 STOCHASTIC OPTIMIZATION AND SEARCH ALGORITHMS

2.2.1 Stochastic Optimization

A stochastic process refers to a random process that captures a random probability distribution model. In mathematical optimization problems, stochastic optimization is a family of methods for minimizing or maximizing an objective function when randomness is present [41]. Stochastic optimization is part of the superset of probabilistic optimization algorithms. Probabilistic algorithms also include robust optimization, distributionally robust optimization, and chance constraint optimization. Heuristics and metaheuristics are important categorizations of approximate methods of mathematical optimization because they refer to how these algorithms perform the approximate optimal solution search and the type of problems they solve. Heuristics refer to informed search methods that systematically explore the search space under a constant heuristic rule. They are problem specific and can be prone to a

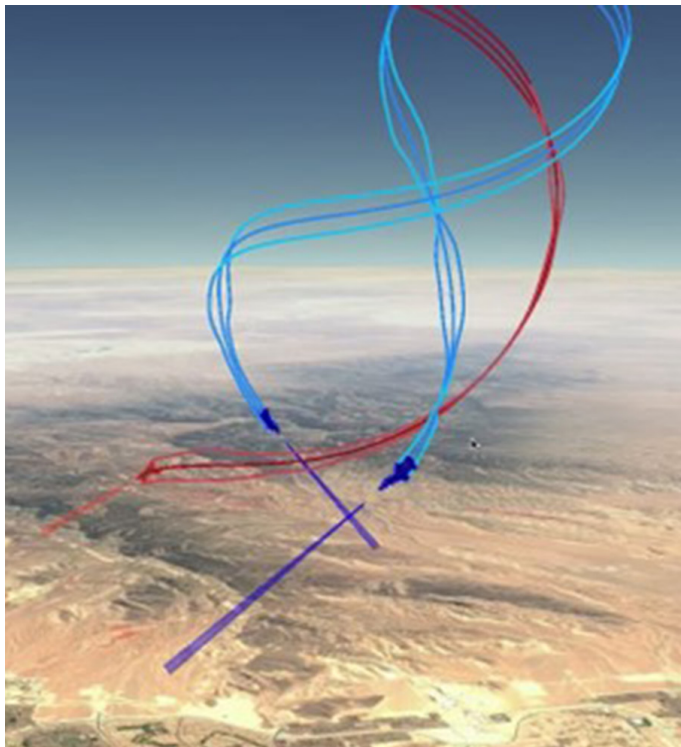


Figure 2-4. DARPA's Air Combat Evolution (ACE) Program Using AlphaDogfight in Simulations (Source: DARPA, <https://www.darpa.mil/news-events/2021-03-18a>).

local optima trap. Metaheuristics are high-level strategies for informing and guiding a search using multiple criteria so that the overall search behavior keeps changing to explore the search space [42]. Metaheuristics are not problem specific [43] and are suitable to multiobjective optimization problems [44].

2.2.1.1 Swarm Intelligence (SI)

SI is the behavior pattern that forms from the decentralized collective actions of distributed, self-organized agents (Figure 2-5). Each individual agent acts and reacts according to its local rules from which a complex group behavior unfolds [45]. The behavior that emerges from the interactions of the agents offers robustness, adaptation, flexibility, and scalability, which are key advantages afforded by the swarm intelligence over actions of an individual agent or those from multiple agents with centralized control [46].

The behavior of the swarm is classified into the following [45]:

- Spatial organization (aggregation, formation, and self-assembly).
- Navigation (collective exploration, coordinated motion, and swarm transport).
- Decision making (consensus, task allocation, group perception, and collective fault detection).
- Other paradigms like human-swarm interaction.

Although SI algorithms can suffer from scalability issues, they have been successfully applied to solve several optimizations, clustering, planning/routing, scheduling/load balancing, and collision avoidance problems. (See Section 3.4.2 in this report for select state-of-the-art examples of swarm intelligence algorithms and their use-cases in autonomous systems behavior and path planning.)

2.2.1.2 Evolutionary Algorithms (EAs)

EAs are nature-inspired, stochastic search algorithms that employ the principles of evolution—reproduction, genetic crossover, and mutation to improve the outcome of a desired quantity (fitness) in a system [34]. This process leads to an evolved solution set that is better adapted and suited to the environment from which it originates than the global population of potential solutions.

Examples of the traditional EA algorithms are the Genetic Algorithm, Genetic Programming, Evolutionary Programming, Evolutionary Strategies, and Differential Evolution (DE). Covariance Matrix Adaptation-Evolutionary Strategy, Neuroevolution of Augmenting Topologies (NEAT), Memetic Algorithm, and Natural Evolution Strategy represent some of the recent variants of the EA algorithms.

In recent years, EAs were used to improve the performance of AI/ML models in all ML processes—namely, in preprocessing (e.g., feature selection and imbalanced data resampling), learning (e.g.,

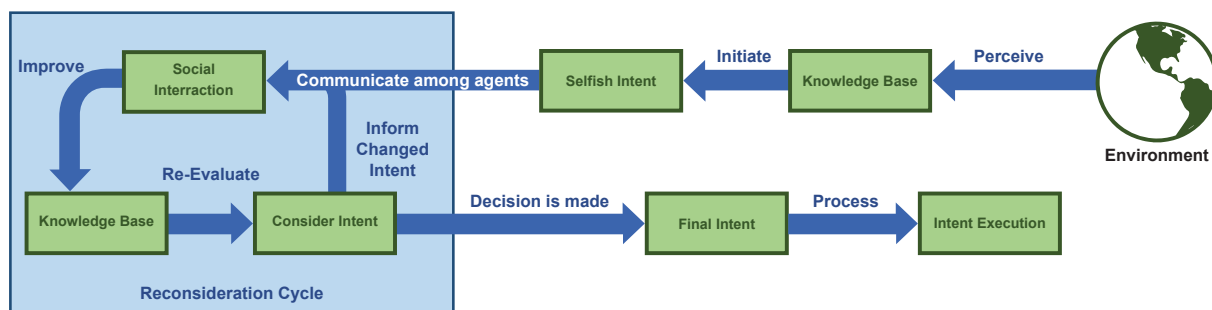


Figure 2-5. Swarm Algorithm Concept for Each Agent (Source: QinetiQ, Adapted From Cunha [47]).

parameter setting, membership functions, and NN topology), and postprocessing (e.g., rule optimization, decision tree/support vectors pruning, and ensemble learning) [48]. EAs have also been successfully used as an effective offline method for multiobjective control parameter optimization in the control-theoretic context [49].

EAs are robust and global compared to traditional gradient-based optimization methods and may be applied without expert knowledge or domain-specific heuristics [34]. Gradient-based optimization methods are prone to finding suboptimal solutions due to their susceptibility to the choice of initial conditions, function accuracy, and search criteria, as the search may be guided to a narrow region and lead to global optima reachability issues. EAs are known to provide global solutions since the entire solution space is sampled from and operated on [49]; however, they do not offer a convergence guarantee toward optima.

2.2.1.3 Physics-Inspired Algorithms

From electromagnetic and gravitational phenomena to material and quantum physics, the following algorithms in this class are inspired by the phenomenon observed in the physical world:

- **Artificial Potential Field Algorithm (APF):** A method that models the environment as an artificial potential field and uses virtual force assignment so that the target point is a gravitational field while points of interest (e.g., obstacles) generate repulsive fields. APF is being used extensively in robotics and aerial path planning. In Xie et al. [50], a modified APF was used for unmanned aerial vehicle (UAV), three-dimensional (3-D) path planning in a threat-infested environment. The method improves the traditional APF local minimum and target nonreachability problems and demonstrates obstacle avoidance in real-time for static and dynamic environment with unknown and moving obstacles. Similarly, Feng et al. [51] showed fixed-wing UAVs

forming around target points while performing obstacle avoidance. In underwater settings, Fan et al. [52] proposed an improved APF method for underwater path planning in an unknown obstacle's environment. While APF is not a stochastic algorithm, it is often combined with stochastic algorithms, such as the genetic algorithm or particle swarm optimization, to further improve its performance [53].

- **Simulated Annealing (SA) Algorithm:** A probabilistic technique in the family of metaheuristic methods for approximating global optimum of a function over an arbitrary (discrete or continuous) search space replete with local optima. Inspired by the physics of materials and their properties during heating and cooling phases, SA can solve optimization problem whose objective functions and constraints are multivariate and used for local and global searches.
- **Quantum-Bacterial Swarm Optimization (QBSO) Algorithm:** A semiphysics-inspired algorithm incorporating bacterial foraging swarm behavior with quantum theory. This algorithm is an improved version of the Bacterial-Foraging Optimization (BFO) algorithm, as the latter does not solve discrete problems. By introducing the Quantum Effect, the QBSO adapts the foraging behavior of the BFO to accelerate the convergence rate [53].

2.2.1.4 Other Metaheuristics

Additional optimization methods within the metaheuristics family of algorithms are those advocating integrating adaptation or learning search heuristics to intelligently escape local minima or arrive at global optima, such as hybrid optimization methods and reactive search methods. They include the following:

- **Guided Local Search (GLS) Algorithm:** A version of the local search algorithm with the aim to improve efficiency and robustness. GLS is a penalty-based, metaheuristic algorithm. The key improvement in GLS is its guidance to

escape from local optima solutions and find better solutions via its use of a penalizing mechanism that determines which features are selected to penalize when the local trap occurs. A notable variant of the GLS algorithm, called Elite-Biased GLS (EB-GLS), showed improved performance in vast search spaces of combinatorial optimization problems [54].

- **Tabu Search (TS) Algorithm:** A search method for a local search type of problems in mathematical optimization. One of the main components in TS is its adaptive memory used in the search and responsive exploration. The use of recency and frequency memory in TS fulfills the function of preventing the searching process from cycling endlessly in the search [55]. A noteworthy property of TS is its applicability to combinatorial optimization, where the objective of obtaining an optimal ordered solution applies.

Local search algorithms tend to converge on local regions that may be suboptimal. However, hybrid optimization methods can be effective in several scenarios. For example, EAs are ill suited for fine-tuning parameter structures that are near-optimal solutions. But the main advantage of EAs lies in the quick localization of high-performance regions of a vast and complex search space. When these regions are located, local search heuristics algorithms like TS can be used in conjunction to fine-tune optimal parameters [56].

- **Reactive Search:** A set of methods that merge ML and statistics within a heuristic search to solve complex optimization problems. Reactive search is a learning search through an internal online feedback loop for the self-tuning of critical parameters. Similar to how human brain systematically learns from past experiences, learning on the job, rapid analysis of alternatives, coping with incomplete information, and adaptation to events, the use of ML automates the algorithm selection, adaptation, and integration [2].

2.2.2 Graph Search Algorithms

Graph search algorithms are a family of search algorithms geared toward path-planning applications, such as solving for shortest paths in static and dynamic environments. In this report, we focus on informed search algorithms that are based on the heuristic A* algorithm (an optimized version of the classic Dijkstra's algorithm) and the sample-based search algorithms. These graph search algorithms include the following:

- **D* Algorithm:** An informed incremental graph search algorithm based on the A* for dynamic and unknown environment. The algorithm avoids the computational cost of backtracking and, hence, is faster than the classic A*. It is used to generate a collision-free path in dynamic environment having moving obstacles. The algorithm D* and its variants can be employed for any path cost optimization problem where the path cost changes during the search for the optimal path to the goal, which makes the algorithm fit for online replanning. D* is most efficient when these changes are detected closer to the current node in the search space [57] but is notably more efficient than A*, as it avoids the high computational cost of backtracking.
- **Rapidly Exploring Random Trees (RRT*):** A sampling-based tree search algorithm used to efficiently find the path from a start to an end point in a nonconvex, high-dimensional space with state constraints. RRTs expand by rapidly sampling the space, growing from the starting point, and expanding until the tree is sufficiently close to the goal point. In every iteration, the tree expands to the nearest vertex of the randomly generated vertex. This nearest vertex is selected in terms of a distance metric. It can be Euclidean, Manhattan, or any other distance metric. The algorithm is designed with few heuristics and arbitrary parameters. Many variants of the algorithm offering different sets of improvements exist like those geared toward

applications or computational requirements. RRT* is an optimized version that claims to achieve convergence toward the optimal solution, thus ensuring asymptotic optimality along with probabilistic completeness. As a result, it obtains shorter paths but at the expense of computation performance, as it is slower than the RRT [57]. Notable improved variants for faster convergence are RRT*-smart and informed-RRT*.

While RRT is a faster algorithm than the classic A*, it can produce longer paths than A*. However, both RRT and A-star (A*) are known to suffer from slow convergence speed for 3-D navigation in a dynamic environment (curse of dimensionality) and an additional issue of path smoothness in the case of RRT [50].

2.3 EMERGING AI PARADIGMS

We highlighted vastly applied techniques to the state-of-the-art AI systems. Some techniques like RL can be formulated with knowledge transfer in mind and, hence, can be applied to similar problems. Nevertheless, it remains to a large degree that applications of current paradigms are narrow and lack several key characteristics of strong AI—generalizability, explainability, knowledge abstractions, common sense, and causal reasoning. Current research in neurosymbolic AI and neuroevolution (NE) attempt to address some limitations of narrow AI in different ways.

2.3.1 Neurosymbolic AI

Neurosymbolic AI is an emerging area of AI that combines the classic rules-based AI with modern deep-learning techniques. The architecture emphasizes interaction between neural, symbolic, and probabilistic methods and inference. Symbolic part represents reason with abstract knowledge. The probabilistic inference establishes causal relationship between facts, reason about uncertainty, and unseen scenarios. The neural part discovers representations and patterns to sense

environment data to knowledge and help navigate search spaces [58].

Neurosymbolic models have been shown to outperform state-of-the-art DNN models in image and video reasoning domains. Benchmarking the state-of-art techniques with neurosymbolic ones showed better performance for accuracy and training time than the traditional models [59].

One of the inherent drawbacks in this hybrid modeling technique is the coupled complex control flow compared to traditional NNs, which makes the computation partly unsuitable for parallelism. However, from a computational perspective, neurosymbolic architectures are neural-network centric (parallelizable), which translates to the possibility of separating the symbolic part [59].

2.3.2 NE

NE is the artificial evolution of NNs using genetic algorithms [60]. Stanley and Risto [61], in a foundational publication in 2002, described the NEAT concept. They showed how GA can evolve not just the connection weights, as an alternative to stochastic gradient descent in backpropagation of networks, but both the network structure and connection weights to significantly enhance the performance of the NE network. The algorithm uses the structure to minimize the dimensionality of search space.

With the advent of accessible high-performance and distributed computing, these NE ideas and approaches are reemerging to showcase their true potential. Similar to deep learning, once trained with large datasets and scaled to take advantage of the distributed computing, recent research showed that NE significantly improved the performance of its models. Furthermore, it was found that NE is a competitive alternative to gradient-based methods for training DNNs and reinforcement learning models [61].

SECTION 03

APPLICATIONS OF AI FOR WEAPONS SYSTEMS

AI has the potential for application to many facets of the weapons systems ecosystem. It is used to control systems, thus enabling autonomy and improving the performance to select problems in guidance, navigation, and control in challenging environments. Similarly, AI can be used to solve challenging problems in mission and path planning, thereby achieving a greater level of complex mission objectives and operational requirements. AI is also used in the electronic warfare domain for support, countermeasure, or even counter-countermeasures. It may also be employed in information fusion from various system hierarchies and domains to divulge abstract high-level value battlefield intelligence and provide critical cues and fast-paced decision making that will consequently create a valuable tactical advantage in modern warfare.

This section of the report will highlight the use of state-of-the-art AI methods in the various AI problem domains applicable to autonomous and weapons systems. It is organized according to the following problem domains.

- Autonomy
- AI in perception
- AI in guidance, navigation, and control
- Mission and path planning
- Intelligent strategy
- Opponent modeling
- Cognitive electronic warfare

3.1 AUTONOMY

3.1.1 Definition, Levels, and Frameworks

To present what is meant by autonomy, we state first how it is related to intelligent or robotic agents. Since the context of the discussion is geared toward weapons systems, the terms “robots,” “agent,” and “autonomous system” are used interchangeably when discussing AI.

In a broad sense, intelligence is the quality that enables an entity to function appropriately and with foresight in its environment. The environment in which the agent interacts is essential to understanding and designing an autonomous system. In nature, for example, an agent “or a complex system” is intrinsically adapted to its environment (habitat). Therefore, the level of intelligence needed by the system to perform a task is dictated by the environment involved [62]. The environment is distinguished by its six different dimensions—fully or partially observable, deterministic or stochastic, episodic or sequential, static or dynamic, discrete or continuous, and single agent or multiagent.

The concept of an intelligent system evokes the understanding that at least some functions are intrinsically performed in an autonomous fashion by the system. Otherwise, the questions then become the following:

- What is intelligent about it if it cannot do a nontrivial task without guidance or involvement of a human?

- Is reasoning, learning, or action of the system on the environment autonomous?
- How about if they are all autonomous—is it only then the system is intelligent?
- What level of autonomy makes a system autonomous when deployed?
- Is autonomy simply an automatic system or automated system?

The word “autonomy” is derived from the Greek words “auto” and “nomos,” which mean “self” and “law” or “governing,” respectively. Autonomy applies to a “self-governing” agency, whether it is a human, organization, or cyberphysical system.

There appears to be widely persistent semantic confusion among automatic, automated, and autonomous descriptions of systems. An automatic system is a central component in a cyberphysical system whose function is to perform a specific set of actions with predefined responses (e.g., flight autopilot and a car’s electronic stability control). Automated systems are more complex systems that function with limited to no human operator involvement in a structured environment with expected behavior. The behavior of the automated system is traditionally rule-based, aimed at performing a specific set of tasks according to prescribed rules (e.g., counter rocket, artillery, and mortar [CRAM] weapon).

An autonomous system is characterized by the ability to select and plan an appropriate course of action to reach an objective from its perception of the environment, situational awareness, and understanding of the local or dynamic context [63]. This ability has a direct impact on the amount of monitoring and delegation of tasks made by the human operator to the system. The level, frequency, and nature of interaction with human operators are important aspects in the human-robot interaction (HRI) study.

While there is a large sample of definitions produced by both academia and defense research,

it is challenging and still ill-defined at what point a system becomes autonomous. But with the previous definition in mind, it is reasonable to say that an automated system that is partially adaptive to a changing environment is trailing onto the autonomous realm. In Russel and Norvig [6], a connection was drawn between autonomy and adaptation by learning as follows: “A rational agent should be autonomous. It should learn what it can to compensate for partial or incorrect prior knowledge.”

Often, the confusion is further exacerbated by the assumption that intelligence is an integral part of autonomy. The level of intelligence varies in the different type of agents (reflex agent, model-reflex agent, utility-based agent, goal-based agent, and learning agent), with the learning agent having a higher form of intelligence. In fact, intelligent systems are characterized by their predictive capabilities, uncertainty management, and autonomy. An autonomous system may be intelligent if it learns and understands to solve problems (predictive capability) and make decisions (uncertainty management).

In defense research, the DoD released DoD Directive 3000.09 in 2012 (administratively updated in 2017) as formal policy on autonomy in weapons systems and the DoD AI Ethical Principles in 2021 for the design, development, deployment, and use of AI capabilities [64].

The DoD categorizes a weaponized autonomous system according to the level of autonomy as follows [65] (see Figure 3-1):

- Semiautonomy – a “human-in-the loop” designation. The human operator specifies the target and conditions for engagement. This includes the following:
 - Semiautonomous weapons systems with engagement-related autonomy functions, such as acquisition, tracking, identification, cueing, target prioritization, and timing of when to fire, or providing terminal

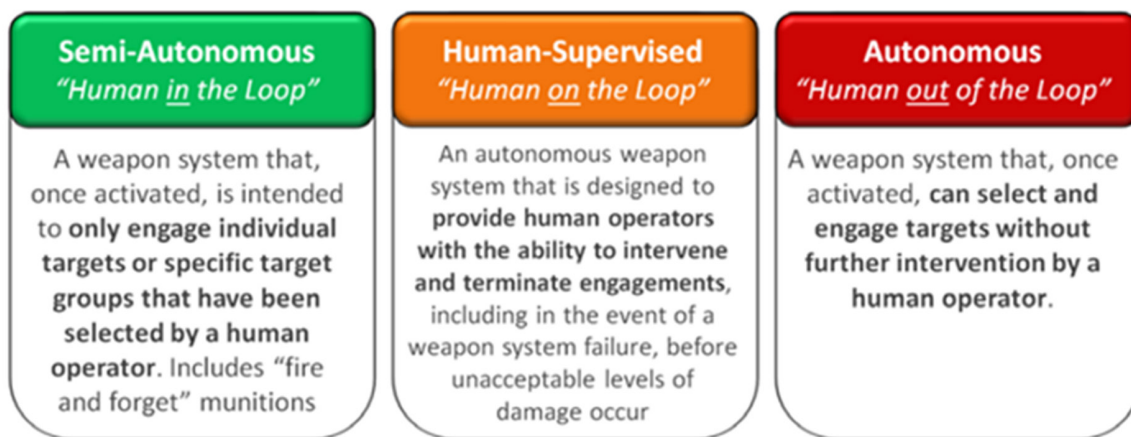


Figure 3-1. DoD 3000.09 Definitions of Top Levels of Autonomy (Source: Deputy Secretary of Defense [64]).

guidance to home in on selected targets, provided that human control is retained over the decision to select individual targets and specific target groups for engagement.

- “Fire and forget” or lock-on-after-launch homing munitions that rely on tactics, techniques, and procedures (TTPs) to maximize the probability that the only targets within the seeker’s acquisition basket when the seeker activates are those individual targets or specific target groups that have been selected by a human operator.
- Supervised Autonomy – a “human-on-the loop” designation. The human operator supervises the weapon’s mission and intercepts engagement in the event of weapons system failure or change in the rules of engagement.
- Full Autonomy – a “human-out-the loop” designation. The weapons system initiates its target selection and proceeds with engagement without the intervention of the human operator.

It should be noted that all autonomous systems are supervised by a human at some autonomy level and the autonomous systems’ software operates with the encoded designed limits on the actions and decisions [64].

Several categorizations of autonomy levels exist, the most prominent of which attempted to provide primitive models for autonomous behavior in terms of the level of human-robot interaction (HRI) independently of context. Some of these models are “Sense, Plan, Act” (SP&A) (see Table 3-1 [66]), “Think Look, Talk, Move, and Work,” and “Observe, Orient, Decide, and Act.”

To facilitate the characterization and articulation of autonomy in unmanned systems, the National Institute of Standards and Technology (NIST) devised a domain-agnostic framework called the “Autonomous Level for Unmanned Systems” (ALFUS) framework [67]. The DoD’s Defense Science Board introduced a domain specific framework named “Autonomous Systems Reference Framework” (ASRF) [68].

The ASRF taskforce posited that defining taxonomies and grouping functions based by discrete levels of autonomy emphasize the focus on the robot rather than the collaboration between human and robot to achieve the desired capabilities, results, and effects. They recommended abandoning the use of “levels of autonomy” and replacing them with the following:

- An emphasis on the explicit allocation of cognitive functions and responsibilities between human and robots when making design decisions.

Table 3-1. SP&A-Based Levels of Autonomy [66]

Level	Sense	Plan	Act	Description
1 – Manual	H	H	H	The human performs all aspects of the task, including sensing the environment, generating plans/options/goals, and implementing processes.
2 – Teleoperation	H/R	H	H/R	The robot assists the human with action implementation. However, sensing and planning is allocated to the human. For example, a human may teleoperate a robot, but the human may choose to prompt the robot to assist with some aspects of a task (e.g., gripping objects).
3 – Assisted Teleoperation	H/R	H	H/R	The human assists with all aspects of the task. However, the robot senses the environment and chooses to intervene with the task. For example, if the user navigates the robot too close to an obstacle, the robot will automatically steer to avoid collision.
4 – Batch Processing	H/R	H	R	Both the human and robot monitor and sense the environment. The human, however, determines the goals and plans of the task. The robot then implements the task.
5 – Decision Support	H/R	H/R	R	Both the human and robot sense the environment and generate a task plan. However, the human chooses the task plan and commands the robot to implement actions.
6 – Shared Control With Human Initiative	H/R	H/R	R	The robot autonomously senses the environment, develops plans and goals, and implements actions. However, the human monitors the robot's progress and may intervene and influence the robot with new goals and plans if the robot is having difficulty.
7 – Shared Control With Robot Initiative	H/R	H/R	R	The robot performs all aspects of the task (sense, plan, and act). If the robot encounters difficulty, it can prompt the human for assistance in setting new goals and plans.
8 – Executive Control	R	H/R	R	The human may give an abstract high-level goal (e.g., navigate in environment to a specified location). The robot autonomously senses the environment, sets the plan, and implements an action.
9 – Supervisory Control	H/R	R	R	The robot performs all aspects of the task, but the human continuously monitors the robot, environment, and task. The human has override capability and may set a new goal and plan. In this case, the autonomy would shift to executive control, shared control, or decision support.
10 – Full Autonomy	R	R	R	The robot performs all aspects of a task autonomously without human intervention with sensing, planning, or implementing action.

Note: H = human and R = robot.

- Tailoring these allocations based on mission phase and echelon.
- Making high-level system trades in the design of autonomous capabilities visible.

3.1.2 Autonomous System's Functional Components

We describe the general concept of an autonomous system's main software functional components for the purpose of mapping state-of-the-art AI techniques to AI problem domains in autonomy (Figure 3-2). Typical robotic autonomy core's top-level functional components include perception pipeline, planning, and vehicle control.

3.1.2.1 Perception Pipeline

Perception provides the ability to sense and transform raw sensorial inputs (proprioceptive, exteroceptive, and abstract) into usable information via the capture, representation, and interpretation of environmental cues (e.g., location, geometry, motion, spectral content, etc.) for the purpose of mission planning, motion control, countermeasures, and information intelligence [69]. In autonomous systems, perception plays a central role in the following [63]:

- Guidance, Navigation, and Control: Perception supports the path planning, dynamic

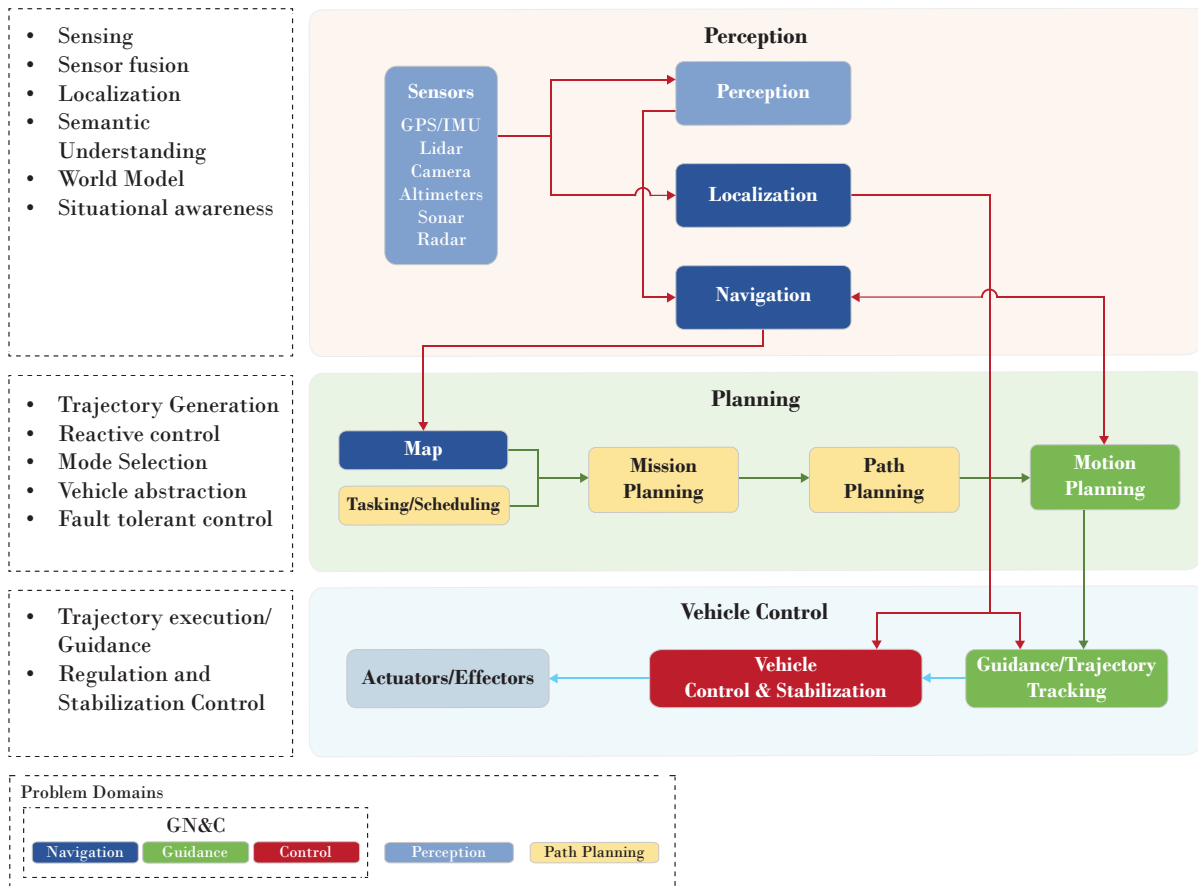


Figure 3-2. Autonomy Functional Components (Source: QinetiQ).

- replanning, motion control for single-agent systems, and motion coordination in multiagent systems.
- **Planning:** Perception provides mission sensing and abstract cues for strategizing mission and scenario planning, teaming, and coordination.
- **Diagnostics and Fault Management:** Perception provides fault detection and system health management by detecting anomalies and fault patterns leading to failure for subsequent online mitigation or management (functional component not explicitly shown).
- **Situational Awareness:** Perception provides the hierarchical information about targets and obstacles in the environment to detect and articulate context understanding in scenes and scenarios. Perception is key in obstacles detection and discrimination (e.g., detection and discrimination between submerged

moving targets, mines, and underwater dwellers). Perception may be externally aided by additional data from distributed perceptive sources via ground station or command and control systems.

Sensors vary widely in autonomous systems, depending on the system mobility platform and the environment. These devices can provide valuable information about the state of the system (proprioceptive sensors) as well as changes made to the environment or experienced by the system as a result of its interaction with the environment (exteroceptive). Examples of exteroceptive sensors are Global Positioning System (GPS), magnetic field sensor, sonar, radar, altimeter, laser rangefinder, red-green-blue (RGB) camera, and spectral sensors. Examples of proprioceptive sensors include inertial measurement unit (IMU), alpha/beta vanes, anemometers, and odometers.

In perception, for example, sensor fusion is used to combine visual target state data from cameras, ranging devices such as light detection and ranging (LiDAR) or sonar, and velocity from Doppler radar to provide more accurate positioning and tracking information of targets. Information fusion may also be used to incorporate abstract information (e.g., temporal cues of activities, profile, image registrations, and georegistration) with specific models to obtain high-level abstraction object state estimates.

While it may be external to perception pipeline, the localization function provides continuous estimates of the autonomous system position, velocity, attitude, altitude rate, and acceleration among other states pertaining to platform motion state in its environment. The localization layer in the autonomy stack leverages typical domain-specific sensor fusion techniques suited for specific mobility and environment.

The navigation function is responsible for generating a map of the autonomous system's local environment and the detection of any hazards that might impede mission progress so that a collision-free navigation path from its current location is determined. To execute the mission, this function processes data from the localization and perception functions (sensing the environment) to ascertain the autonomous system's behavior or motion on the map.

3.1.2.2 Planning

The planning function produces a sequence of actions or behaviors the autonomous system must follow from a specified starting configuration (or position) to its final goal (or destination) while avoiding obstacles and appearing impediments in its path in accordance with mission objectives. These mission objectives may be metric (e.g., smooth trajectories) or symbolic (e.g., avoid georegistered obstacle). From the perception layer, the planning function receives the autonomous system position, detected obstacles information,

map, and mission objectives to solve three layers of planning abstraction—mission planning, path planning, and motion planning. State-of-the-art planning algorithms feature adaptive capability to the changing environment, such as agile mission replanning to engage targets of opportunity. This feature requires sensor cuing, context change detection, root cause analysis, mission impact analysis, and agile replanning.

3.1.2.3 Vehicle Control

In the control function, the combined outputs from navigation, planning, and perception are transformed into allocated commands for the different effectors on the mobility platform, therefore affecting its motion. Platform motion is controlled by different levels of control algorithms. Medium-level controls (e.g., autopilot) are either the guidance algorithms that ensure the vehicle trajectory closes in on the target trajectory or trajectory tracking algorithms that ensure the tracking of the generated trajectories in the planning layer and that the path deviation error is minimized. Low-level control are algorithms that stabilize and regulate motion via actuations. The control function can also have an aggregate control function, such as payload control or weapons subsystems control.

3.2 AI IN PERCEPTION

We highlight the state-of-the-art technologies associated with perception in the typical autonomy top-level functional architecture. These AI techniques apply to the perception pipeline in both autonomy core software and payload ISR systems. AI techniques for autonomy include techniques to support the understanding of the environment and support for motion and behavior control, such as obstacle detection, self-localization and mapping, and motion state estimation. AI techniques for ISR systems support situational awareness, scene understanding, automatic target, and activity recognition and tracking.

3.2.1 Image Segmentation

Semantic understanding is the task of associating semantic meaning to image content to construct understanding and context. Image segmentation is particularly important to scene or visual understanding. There are three main categories of segmentation—semantic, instance, and panoptic (see Figure 3-3).

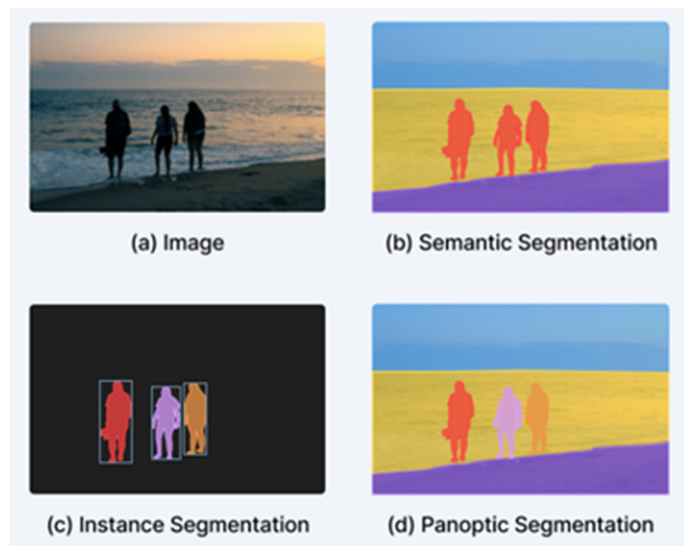


Figure 3-3. Semantic (b) vs. Instance (c) vs. Panoptic (d) Segmentations (Source: V7Labs).

3.2.1.1 Semantic Segmentation

Semantic segmentation aims to classify objects in an image at the pixel level by the following three subtasks [70]:

1. Classifying a certain object in the image.
2. Localizing it by finding the object in a bounding box.
3. Grouping the pixels in a localized image by creating a segmentation mask.

Important improvements have been achieved in recent years, and some of the prominent CNN architectures for semantic segmentation are DeepLabV2, PSPNet, and ParseNet.

3.2.1.2 Instance Segmentation

Instance segmentation refers to the technique of detecting, segmenting, and classifying every individual object in an image. This technique combines semantic segmentation and object detection and therefore produces a richer output format as compared to object detection and semantic segmentation networks separately [71].

A state-of-the-art model for instance segmentation is Mask Region-based Convolutional Neural Network (R-CNN). Based on the ResNet topology (ResNet-50), this augmented network provides object segmentation with added mask information. The outputs are class label of the detected object (ResNet-50), bounding box offset (feature pyramid network), and an output of the object mask for the purpose of extracting a much finer spatial layout of an object [71].

3.2.1.3 Panoptic Segmentation

Panoptic segmentation is a state-of-the-art image segmentation technique that unifies semantic and instance segmentation methods. Objects to be segmented are assigned two labels to each of the pixels of an image: (1) semantic label and (2) instance ID. State-of-the-art panoptic segmentation is the DETection TRansformer (DETR) framework that provides end-to-end object detection. The framework showed accuracy and run-time performance on par with the highly optimized, faster R-CNN baseline and outperformed competitive baselines on the Common Objects in COntext (COCO) object detection dataset [72]. Another common related framework is the Detectron2.

3.2.2 Target Detection, Classification, and Scene Understanding

In the next sections, we highlight notable advances in perception in major image-based database challenges using state-of-the-art DNN topologies for image and video data.

3.2.2.1 Image Techniques

The following techniques are intrinsic perception technologies of visual knowledge extraction, particularly in object classification, image context, and human-machine interaction [3]:

- **Image Classification:** The ability to categorize what is seen in the image. Image recognition is used to identify pedestrians in car driving, categorize specific objects in an image or video, detect tumors in diagnostic system, and detect defects in productions. In the 2021 ImageNet Challenge, Top-5 (five first probabilities aligned with the labeled image) was achieved by Microsoft AI's Florence-CoSwim-H model, with a 99.02% accuracy.
- **Image Generation:** The ability to generate images that are indistinguishable from real ones. This task typically uses GANs, and the similarity metric (similar to real images) uses the Frechet Inception Distance score.
- **Human Pose Estimation:** The ability to estimate different human body posture and joint kinematics to correctly label the human pose for different applications—activity analytics, target surveillance, and crowd activity monitoring. The latest advances in deep learning models achieved 99.5% human pose estimation accuracy on the Leeds Sports Poses dataset.
- **Semantic Segmentation:** The ability to assign individual image pixels a class or category, such as a human, background, or building. Autonomous systems applications, such as a self-driving car, require pixel-level image segmentation (to identify which parts are a human in the field of view and which are other cars and roads), image analysis (distinction between background and foreground in photos), and aspects that deal with occlusions. The 2021 top-performing AI systems on the challenging Cityscapes dataset (database containing images from urban street environments from 50 cities taken during the

daytime in different seasons) achieved 86.20% accuracy.

- **Visual Reasoning:** The goal of this task is to develop an AI that can reason broadly across a combination of visual and textual data. While existing AI already excel at narrow visual classification, face detection, and object segmentation at a level far exceeding humans, this is a challenging and active area of research because it requires more abstract reasoning to generate valid inference about actions and intent of subjects in an image.
- **Visual Questioning Answering (VQA):** This area combines language understanding, vision, and common-sense reason in an AI. This technology answers some of the challenges in the Explainable AI domain in which the AI answers open-ended questions about images at a high level. Top-performing AI in the 2021 VQA challenge achieved 79.8% accuracy close to a human baseline of 80.8%. An AI created by a collaboration between Google Research, Michigan State University, and Brown University achieved 89.6% accuracy on the Kinetics-600 dataset (database of video displaying a wide range of human activities).

3.2.2.2 Video Techniques

The listed video techniques use the temporal information between image frames to detect patterns and infer visual knowledge central to ISR and autonomous weapons, such as target acquisition and cueing, situational awareness, scene understanding, and human-machine interaction [3].

- **Activity Recognition:** The ability to identify activities that occur in videos ranging from walking to more complex, coupled activities like preparing for something good or bad. This requires the AI to recognize and chain discrete actions together to determine the abstract top-level activity.

- **Object Detection:** The ability to identify objects within an image. This involves two major set methods: (1) one that prioritizes speed for real-time operation but is slightly less accurate—You Only Look Once (YOLO), RetinaNet, and SSD and (2) two-stage methods that prioritize accuracy, such as Faster R-CNN, Cascade R-CNN, and Mask R-CNN. YOLO achieved a mean average precision (MAP) of 80.7% by 2021, while Faster R-CNN achieved 87.69% MAP.
- **Visual Commonsense Reasoning (VCR):** A new benchmark for visual understanding aimed at answering questions about a scenario presented from image frames and providing reasoning behind the answers. Since the introduction of this challenge, visual commonsense reasoning has improved significantly; however, the scores are still below the human baseline of 85%. The current best mark on VCR is about 72%. Improvements have become increasingly marginal, which suggests new techniques may be required to further improve performance closer to human baseline.

3.2.3 Sensor Fusion

Sensor fusion concerns the ingestion and unification of information from a variety of sensors taking different measurements. Multisensor data processing goes back to the early inertial navigation systems when signals from gyroscope signals, accelerometers, and magnetometers were combined to create a less uncertain state estimation than the separate measurements from each different sensor. In the context of an autonomous system, multisensor fusion is the exploitation capability to decode information from multiple heterogenous sensors to infer context and obtain a fused singular “view” of the sensed environment. AI techniques have been leveraged to improve this important processing function in perception and localization. Examples of sensor fusion of perception are described next.

Blasch and Zheng [73] presented a multimodal imaging fusion for the simultaneous context-

aided tracking and identification of physics-based and human-derived information. The modalities exploited were electro-optical visual imaging and infrared imaging. The fusion augmented the imagery with context via content colorization and enables environment interpretability for explainable intelligence.

Huang et al. [74] proposed a novel integrated DNN structure that leverages multimodal sensor fusion processing and scene understanding for end-to-end autonomous driving. The structure provides concurrent scene segmentation and guidance commands to vehicle steering and speed control systems. The structure is composed of three networks—multimodal sensor fusion encoder, scene understand decoder, and conditional driving policy network. Multimodal sensor fusion encoder is the ResNet-50 V2 structured residual NN that receives the concatenated RGB and (LiDAR) depth imagery and outputs the features map. The scene understanding decoder processes the features map via deconvolutional layers, with a softmax activation function for the last layer and ReLU activation functions for the previous layers. The output of the scene understanding is the categorization of each pixel in the original image, with explicit expressions of the driving scene. The conditional driving policy network gets the global average pooled feature map as an input from the multimodal sensor fusion encoder and outputs the desired speed and steering control signals. When the structure is trained as a whole, it shows 100% success rate in navigation tasks in training and unobserved circumstances.

3.3 AI IN GUIDANCE, NAVIGATION, AND CONTROL

Control is a broad term but generally used within a specific context. Consider the context of multidomain battlefield management. Control in this context may refer to the tactical instructions generated by the decision making of high-echelon commands or the control of operations, such as logistics, inventory control, and resource

allocations. In command and control, control may refer to the supervisory task of systems and humans, but it may also refer to instantiation of mission and online planning. In multiagent systems, control is specified in terms of how agents are controlled, in which case, it can be centralized, decentralized, or hierarchical. In its low-level representation, control addresses low-level tasks, such as tracking and regulation functions found in the guidance, navigation, and control (GN&C) problem domain.

We provided an overview of the research and application of AI in the traditional field of GN&C. While the relation of GN&C to the more modern framework of autonomy was articulated in Section 3.1.2, GN&C functions were not explicitly defined. We first explained what GN&C functions solve in an autonomy stack. We then made a contrast between the traditional GN&C methods in conventional weapons and those aided by AI for improved systems performance, followed by the GN&C areas in which AI techniques (navigation and localization, intelligent control, and system identification) were integrated.

3.3.1 GN&C Systems

GN&C are important problem domains that address the design of system segments controlling and supporting the missions of bodies in motion. From putting satellites in orbits and guided munitions to navigating submarines, these systems are absolutely necessary to autonomous systems' motion control and subsequently the achievement of mission objectives. GN&C systems are described as follows:

- A navigation system that measures the instantaneous state of the vehicle. The state of the vehicle contains kinematic states, such as position, velocity, attitude/pose, and uncertainties in the sensor measurements, and, if applicable, other states particular to the vehicle mobility, such as angle-of-attack angle and side-slip angle.

- A guidance system that computes the optimal trajectory and the corresponding vehicle steering commands so optimal trajectory is realized. The guidance system continuously recomputes the remaining path and desired altitude of the vehicle to achieve the mission goal. Examples of guidance algorithms are Line-of-Sight, Pure Pursuit, and Constant Bearing.
- Control system receives the steering commands from the guidance system and steers the vehicle to follow the desired altitude in the presence of all disturbances. The control system has three major functions:
 1. Stabilize the vehicle throughout the mission.
 2. Steer the vehicle to follow the desired altitude as dictated by the guidance system.
 3. Maintain vehicle loading within desired limits.

3.3.2 Conventional Control-Theoretic Methods

The conventional process of designing an automatic control system entails: (1) physics-based modeling of the system dynamics and environment typically described by differential equations and (2) a control architecture suitable for the system case and performance requirements at hand. Estimation of errors (process parameters, uncertainties, and sensor errors) are also part of the modeling effort. For systems for which there is no explicit mathematical model, system identification techniques are used to build an approximate model that describes the system from the input and output data.

While control theory-based methods provide numerous, rigorous frameworks (robust, optimal, nonlinear, adaptive control frameworks, etc.) for the analysis and controller synthesis of large class of control problems centered around global stability guarantee, adapting to complex nonlinear systems in a time-varying environment, however, remains a challenge and constitutes a major area

of research in the control-theoretic paradigm. Model representations of disturbances, errors, and systems parameters variation are far from exhaustive or fully defined to globally represent the uncertainties in the environment. As such, changes in the environment may introduce uncertainties that are unknown and, hence, may not be characterized during controller design.

Consider the objective of designing a control system that allows for extreme maneuverability of a UAV (e.g., aerial acrobatics and missile evasion) or a humanoid robot walking over a wide range of terrains and obstacles. While physics-based modeling of the system (flight mechanics, anthropomorphic robot kinematics/dynamics, actuators, etc.) is straightforward, the inherent nonlinearities and complex large environment state space in this case render the control design using control-theoretic approach a difficult undertaking. Furthermore, an accurate closed-form mathematical expression of the performance criterion that captures implicit systems dynamics, environment interaction, and control limitations is not easily derived [33, 74].

A control-theoretic approach like model reference adaptive control is the closest to a more generalized control scheme that ensures stability during system parameter variation, unmodeled disturbances, or environment changes. But the stability guarantee assumes a linear or linearized model of the system. Therefore, if the model is not matched to the actual system, then the closed loop stability is no longer guaranteed [75]. Furthermore, adaptive control schemes suffer from extensibility limitations and are often prone to numerical tractability issues [34]. In terms of nonlinear control, geometric methods, such as adaptive sliding mode control, are prone to the chattering effect [76]. Additionally, solving a discrete optimal control problem by dynamic programming may lead to intractable solutions due to the large number of states and, consequently, the curse of dimensionality [77].

The goal of intelligent control is to provide an alternative approach for adaptive nonlinear control of complex dynamics. An advantage of intelligent control is the built-in ability to learn and adapt to unknown dynamics in the environment. An example of such adaptation is flight in wind-varying conditions, a robotic manipulator arm handling a different object, and a quadruped robot walking over changing terrains.

3.3.3 Intelligent Control

Intelligent control refers to advanced techniques that combine ML with theory-based techniques for the control of complex, nonlinear, dynamic systems. Being an active area of research, common names in the academic research are data-driven control or ML control. In this section, we discuss AI-based techniques applied to control—hybrid methods (DNN + control-theoretic scheme and DeepMPC), reinforcement learning, and imitation learning.

3.3.3.1 DNNs

As stated in Section 2.1.1, DNNs can be used to approximate highly complex multivariate mapping; when combined with adaptive control schemes, hybrid methods provide improved online adaptation to time-varying dynamics [78].

An interesting control strategy/scheme from Shi et al. [79] demonstrated precise quadrotor landing by integrating deep-learned dynamics with a nonlinear, discrete, fixed point iteration controller. Prior related contributions provided no rigorous treatment of the effect of learning on the stability. The significance of this noteworthy approach stems from the exponential global stability guarantee provided that the training errors are bounded. This is achieved via formulating the training error cost function with Lipschitz-bound and spectrally normalized layers of the DNN. The network used ReLU activation functions to eliminate the possibility of a vanishing gradient associated with other activation functions; it learned coupled

unsteady aerodynamics and vehicle dynamics, which were rolled into the synthesized nonaffine control input [80].

In Lin et al. [81], an intelligent entry guidance algorithm was proposed that has a variation of the Numerical Predictor-Corrector Guidance (NPCG). The NPCG determines the bank angle amplitude and direction based on the prediction of a trained DNN and constraint management algorithms (path and terminal constraints). The DNN was trained to predict the nonlinear mapping between flight states and cross/down ranges in real-time. The algorithm corrected longitudinal trajectory and addressed the lateral bank reversal.

The work demonstrates relative fault tolerant predictions such that if the DNN prediction output is outside acceptable error bounds, the dynamic propagation technique takes over the range prediction. The DNN-based prediction offers significant improvement of real-time performance while meeting the constraints requirement. This trained DNN demonstrates operational utility of a DNN model for intelligent system control of a hypersonic munition during reentry maneuvers. In simulation tests, it performs as well as traditional state-of-the-art predictor-corrector algorithms currently in use by most systems.

Du et al. [82] presented a Bank-To-Turn lateral control scheme for reentry of a hypersonic vehicle that combines a control-theoretic approach, namely nonlinear generalized predictive control (NGPC) with a proposed Self-Organizing Recurrent Function Link Network (SORFLN). The RNN-based SORFLN adaptively corrects for the errors in the nominal NGPC during disturbances and uncertainties and dynamically reorganizes to reduce the number of learning parameters. Simulation results showed that the proposed algorithm overcame large uncertainties or disturbances while showing lateral-maneuver stability; this effectively demonstrated utility of a small size DNN for satisfactory performance in reentry control.

Takahashi et al. [83] proposed an adaptive compensator for the accurate trajectory control of a three-link robot manipulator using quaternion RNN. The network was trained by feedback error, and the RNN's adaptive compensation input is synthesized online and added to the control input as a torque command to the manipulator. Test results showed that the proposed method was an effective, adaptive control scheme, once again showing the applicability of AI systems for use in control software.

3.3.3.2 Reinforcement Learning

Reinforcement learning was presented in Section 2.1.2 as a family of learning algorithms suitable for the action control of stochastic processes. It is also applied to high-level abstract motion control of robotics and autonomous systems. We present select applications of RL-based methods to motion control.

Kumar et al. [84] presented a model-free, RL-based algorithm for the rapid motor adaptation of a quadruped robot. The algorithm executes online adaptive locomotion control via base policy network and an adaptation network. The base policy network and adaptation network are trained offline in a simulated environment and then deployed on the robot platform to run asynchronously. The training takes place in the physics-based simulation in two phases. The first phase trains from past actions as well as the output from an encoder network that maps environment parameters (e.g., mass, friction coefficients, terrain height, and motor constants) to a latent extrinsics vector. The base policy network and encoder are trained jointly in the first phase. In the second phase, the adaptation network trains from the history of states and actions on the on-policy data to predict the extrinsics vector. The significance of the method resides in the adaptability of this outer loop, RL-based control scheme to handle a plurality of environmental changes, such as terrain type and composition, payloads, wear and tear, and other unseen disturbances.

In Dooraki and Lee [85], a bioinspired flight controller of a quadcopter based on RL was demonstrated in a simulated environment against a conventional, model-based control scheme called Model-Predictive Control. The RL used a modified Proximal Policy Optimization method to train the controller autonomously—from a tabula rasa (clean slate). The results demonstrated fast adaptation and better maneuvering capability.

Chithapuram et al. [86] used RL's Q-learning algorithm for UAV guidance toward a target and analyzed the results against the Proportional Navigation Guidance (PNG) algorithm. It was found that the guidance law based on Q-learning performed better compared to PNG with a 62.67% hit rate compared to the 18.67% hit rate of the PNG.

In Wang et al. [87], a DRL-based backstepping controller was used to control an air-breathing hypersonic vehicle with actuator constraints (magnitude and rate constraints). A Lyapunov function was selected for the nominal backstepping controller, and a DRL-based parameter adjustment algorithm was used to tackle the controller performance degradation due to the system actuator's constraints and aerodynamic coefficients' uncertainties. Monte-Carlo simulation showed tracking stability and robustness of the proposed schemes against 20% model parameters variations from nominal values.

Furfaro et al. [88] used extreme learning machines based on the RL algorithm Advantage Actor Critic for the guidance and control of hypersonic reentry vehicle under maximum heat-rate avoidance and control constraints.

Iwasaki and Okuyama [89] developed an intelligent control stage based on DRL that extends the capability of an existing control theoretic scheme and improves its performance. The extension adds arbitrary control that would otherwise be impossible to design via control-theoretic approach. The example provided is an inverted pendulum command from swing to inverted

stabilization using an intelligent reference signal for each state of the system that combines with the reference signal of the existing controller.

3.3.3.3 Imitation Learning

In cases when the reward function is sparse or unobtainable due to the complexity of the task, imitation learning offers an optimal policy by imitating the expert decision and actions. Imitation learning is useful when an expert's desired behavior demonstration is easier than specifying a reward function.

In Abbeell [90], the flight controller of an autonomous helicopter was trained by apprenticeship learning, a form of imitation learning, using supervisory input of pilot commands and maneuvering to guide the reinforcement learning process. The results demonstrated that the flight control not only mimicked the pilots' maneuvers but performed exceedingly difficult maneuvers that were challenging to attain via conventional controller design due to the highly nonlinear nature of the dynamics.

Generative adversarial imitation learning (GAIL) is a form of imitation learning that can train policies without explicit definition of a reward function. In Couto and Antonelo [91], a GAIL-based solution was proposed for autonomous navigation of a vehicle in the realistic CARLA simulation environment for urban scenarios. Two similar GAIL architectures were compared where one was augmented by behavioral cloning. The simulated tests showed that both architectures demonstrated expert trajectory imitation capabilities from start to end of the trajectory. However, the behavioral-cloning, augmented network achieved better results in terms of convergence time and training stability.

GAIL has also been used to control the locomotion of the MuJoCo humanoid skeleton using 3-D pose estimation from a motion capture video of human walkers as expert input [92].

3.3.3.4 Deep Model Predictive Control

Model predictive control (MPC) is a control-theoretic scheme where a model is used to predict the future behavior of the system over a finite horizon (time interval). The optimal control input, given a control objective and subject to system constraints, is computed based on the model output (predictions) and the current estimation or measurement of the state [93]. Deep model prediction control (DeepMPC) is an approach to model learning for predictive control designed to handle variations in a system's (e.g., robot) environment and variations during system actions. As the name suggests, the approach is constituted of a DNN structure combined with the MPC scheme [94].

Bieker et al. [95] presented a novel DeepMPC framework in aerodynamics and fluid flow applications. The framework used only low-rank features of the fluid flow to improve control performance considerably. RNNs were used to predict the dynamics of the reduced order of the fluid state. The framework was tested against several computational fluid simulation models (finite volume numerical solver for the incompressible 2-D Navier-Stokes equations). The results indicated good performance and showed the effectiveness and potential of DeepMPC for the control of high-dimensional systems.

When considering the potential use of DeepMPC in controlling a robot, one important consideration is the objective function, which must be precisely defined by a human engineer. In layman's terms, this is similar to programming a robot to act with a specific goal in mind, such as getting from A to B as quickly as possible without an unacceptable risk of crashing. In the fluid flow study, the objective functions related to stability of fluid flow around simulated rotating cylinders. The study aimed to show how the cylinders could rotate without destabilizing the fluid flow.

A survey of the literature on AI methods for the low-level controls of dynamics system shows that current techniques (e.g., reinforcement learning and deep learning) do not provide a global stability guarantee. Hybrid methods like DeepMPC or DNN+ adaptive controller offer limited stability guarantee. As of the writing of this report, AI techniques like reinforcement learning and DeepMPC are mostly applied to high-level control (outermost loops). The control-theoretic approaches are applied in low-level controls in autonomous systems to guarantee stability and meet safety-critical requirements.

3.3.4 Localization and Navigation

In robotics or ground autonomous system applications, localization provides simultaneous positioning and mapping within the autonomous system's environment. The perception layer may provide visual odometry, objects tracking, and points cloud data within the environment that are then used by the localization layer to generate grid maps, terrain data, and positioning of the autonomous system and said objects within the map. Some state-of-the-art localization techniques are simultaneous localization and mapping (SLAM)-based algorithms, such as ORB-SLAM [96] and Dense Piecewise Planar Tracking and Mapping [97].

In UASs, localization produces pose estimate in the inertial frame. This is achieved via the state estimator fusing data from a plurality of sensors (e.g., IMU, GPS, Kollsman altimeter, Doppler altimeter, visual odometry, digital compass, etc.). The techniques used are the unscented Kalman filtering, extended Kalman filtering, particle filtering, and square-root-cubature Kalman filtering.

In underwater unmanned vehicles (UUVs), the localization accuracy is degraded due to the unavailability of global position information in an underwater environment. The accuracy is bounded by periodic resurfacing or the performance of the dead-reckoning algorithms in which error from drift is proportional to the distance traveled.

Accurate navigation is possible using sonar image registration coupled with localization within preloaded underwater acoustic maps of the seafloor or dead reckoning. State-of-the-art techniques use underwater SLAM algorithms [98].

An example of AI-assisted robot navigation algorithm is proposed by Shen et al. [99], in which ML enhances the fusion algorithm (state estimation) in a GPS-denied environment. The algorithm is a hybrid navigation solution called the self-learning, square-root-cubature Kalman Filter that seeks to improve the state estimation and error prediction ability by introducing the LSTM network to provide more stable positioning information during prolonged GPS outages.

3.3.5 System Identification

System identification can be understood as the retrieval of mathematical models that describe system dynamics from active excitation of input and output signals (or output signals only, known as blind system identification). System identification can be expressed as a transfer function, analytical function, differential, or difference equation. Expressing system identification can be a form of art where one must be cognizant of their assumptions about system complexity, model structure, and statistical properties of the underlying dynamics when assessing their results.

The versatility of AI has led to its adoption as a tool for generalized regression and inference problems in science and engineering, specifically, in data-driven modeling of highly complex dynamic systems. Modeling aims only to predict or identify the system under study; it is not concerned with control or recommending action. Weather modeling is an example of system identification. It is aimed at identifying a system to predict the weather; it does not recommend any action. The forecast may inform one's decision to go on a hike tomorrow or not, but the model itself does not make a recommendation.

Examples of the ML techniques applied to system identification and mapping function approximation are described next.

Kube et al. [100] applied autoencoders for learning low-dimensional representation of sample time series of fluctuation-driven flows in Tokamak-confined fusion plasma and for outlier rejection. Sakurada and Yairi [101] showed that autoencoders perform better in anomaly detection when compared to kernel-based PCA in nonlinear dimensionality reduction.

Bakarji et al. [102] applied time-delay embedding and DNN to the model discovery of partially measured nonlinear systems with hidden dynamics and latent or unmeasured variables. The input to the network was the singular value decomposition compressed delay eigenvectors, and the output from the process was reversed to divulge the full-state dynamics. The method was verified against a Lorenz attractor and waterwheel Lorenz system. While the resulting analytical reconstructions did not always match the actual original system, the discovered latent variables and their underlying models share key structural features with the original system "when known." Overall, the method discovers systems models like the original system when initial conditions are perturbed around the original coefficient.

In terms of time-series modeling, Guo et al. [103] proposed an interpretable LSTM network for predicting a multivariable time series signal with exogenous variables. The application of an LSTM network to a multivariable function with exogenous input for the output prediction and time-series forecasting showed the least mean-squared-error (MSE) when compared in a benchmark test to random forest, extreme gradient boosting, and dual-stage, attention-based RNN.

Nonlinear Auto-Regressive with eXogenous Input (NARX) is a powerful tool for system identification and on-line system parameters estimation (Figure 3-4). In Liu and Song [104], NARX was used in




Figure 3-4. MLP-Based NARX Architecture (Source: QinetiQ).

identifying a nonlinear system for control. Its network was trained by a state observer stage (Levenberg-Marquardt algorithm). NARX outperformed several RNN variants, and the results showed the least root MSE.

In recent years, research in AI or ML control methods, namely DRL-based and hybrid methods, such as deepMPC, has showed important results and potential. This is because these methods can tackle several highly nonlinear control problems operating in uncertain and complex environments and address some of the limitations in a control-theoretic paradigm.

As mission complexity and the demand for cross-functional and abstract control increases, control-theoretic approaches (e.g., physics-based maneuvering of combat aircrafts) might become prohibitive due their extensively iterative and laborious design process. These approaches might not meet the adaptability and mission-criticality requirements of human-robot interactions and tomorrow's warfare.

An example of the initiative to use AI in weapons systems control in complex settings is DARPA's ACE program [105] (Figure 3-5).

3.4 MISSION AND PATH PLANNING

Mission and path planning is an important control function in a robotics and autonomy stack. It involves computing the optimal path or quasioptimal path from a source point to a destination point. For dynamic environments, it is a nonpolynomial time (NP-hard) problem whose

complexity increases exponentially with higher degrees of freedom of the states [106]. The optimal path may have to meet multiple objectives under multiple constraints. Path-planning optimization is typically divided into path search methods and trajectory optimization problems [107]. Path search methods are solved by heuristic and metaheuristic algorithms, such as graph methods, swarm intelligence, and evolutionary methods. Path search methods offer a rich set of solutions that attempt to tackle uncertainty and complexity while adhering to the modern control objectives and performance requirements.

Knowledge of the environment determines whether path planning is solved at a global or local level [108]. Global path planning finds the optimal path in a static and completely known environment to the autonomous system. The path is determined from start to final point before the autonomous systems begin controlling its trajectory along the determined path. Local path planning is an imperfect information traversal scenario in which the goal is to find a new path in an unknown or

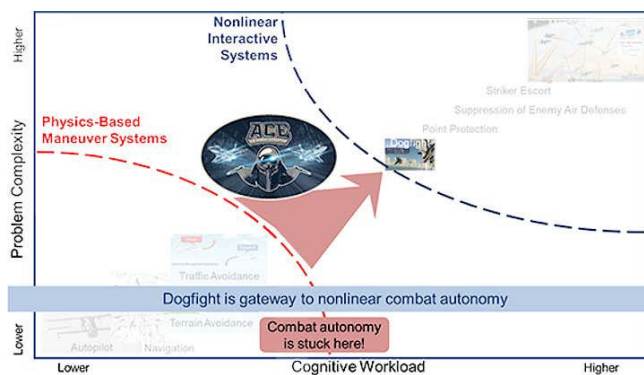


Figure 3-5. Air Combat Evolution (Source: DARPA [105]).

changing environment while the autonomous systems are moving within and sensing its environment.

An effective path-planning algorithm needs to meet the following four performance criteria [109]:

1. Motion planning must provide the optimal path in realistic static environments in a robust fashion.
2. Motion planning must be expandable to dynamic environments.
3. Motion planning must remain compatible with a chosen self-referencing approach.
4. Motion planning must minimize complexity and computational resources.

The following sections survey selected state-of-the-art algorithms applied to mission, path, and behavior planning of single and multiagent systems while meeting these path-planning criteria. Methods briefly summarized include evolutionary-based algorithms like Genetic algorithms (GAs) and Swarm Intelligence algorithms.

3.4.1 GAs

GAs, a subfamily of Evolutionary algorithms inspired by the theory of natural selection in biological evolution, relies on operators, such as mutations or selections in each system. In Kusy et al. [110], a GA was used to control the motion of a swarm of 20 UAVs in a Mobile Ad-hoc Network network topology. Unlike a centralized control with a preplanned mission or rule-based procedure that suffers from their lack of adaptability to a dynamic environment, susceptibility to collisions, and slow operational responses, this near-real-time algorithm is computed by each UAV considering only its state and local information. The UAVs broadcast its global position and position of the neighboring UAVs within communication range and then compute its distance from neighboring UAVs to determine its best motion action [45, 110].

3.4.2 Swarm Intelligence

Swarm intelligence algorithms have been applied to optimization, clustering, routing, scheduling/load balancing, and collision avoidance problems. Similar to evolutionary algorithms, swarm intelligence algorithms have successfully been applied in numerous optimizations, behaviors, and path-planning problems.

One form of swarm intelligence is Particle Swarm Optimization (PSO), which, as a computational optimization method, has been used for optimization, path planning, and parameter tuning of conventional controller schemes. Although the algorithm is known to have a good convergence property to the optimal solution, the fast convergence may lead to premature convergence to a local optimum [111].

Ant Colony Optimization (ACO) is a probabilistic algorithm that models an ant's travel path selection based on the concentration of ants' pheromone deposition (feedback mechanism), which leads to short paths between the colony and food sources. ACO performance depends on the selection of search parameters. Dan et al. [112] proposed a method for optimal selection of parameters for ACO. Suitable for discrete optimization problems, the algorithm is characterized by strong robustness and a parallel computational mechanism [113], even against a noisy sample space (Gaussian Distribution) [114].

In Konatowski and Pawlowski [115], ACO was used for planning of an UAV flight whose path was rerouted to avoid threats. Like PSO, ACO could prematurely converge (suboptimal solutions) in the planning process. Multicolony ACO was suggested to find the optimal solution by a cooperative exchange of information in the search among multiple ant colonies [116].

In swarm robotics, Multiobjective ACO (MOACO) algorithms have showed promising results in constrained multiobjective optimization problems.

Chen et al. [117] employed a modified MOACO to solve task allocation for heterogeneous UAVs by considering target, task, and path-planning objectives. Specifically, the task allocation of heterogeneous UAVs was formulated as a constrained multiobjective optimization problem with three objectives—task benefit, UAV damage, and total range under the physical and operational (logical) constraints. Simulation results showed better performance compared to the standard MOACO and Multiobjective PSO.

Another swarm intelligence algorithm relevant to swarm robotics is the Wolf Pack Algorithm (WPA), which was inspired by the intelligent collaborative wolf behavior during hunting. In Wu and Zhang [118], a search-attack mission planning for UAVs based on the WPA was solved by first computing the optimal solution to the multitarget search problem in an unknown environment modeled after the wolf scouting behavior using the cooperative search algorithm. Then, a distributed self-organizing task allocation algorithm was used for the swarm cooperative target strike, modeled after wolves' labor division behavior. The algorithm has good robustness and global convergence property and is less susceptible to the premature convergence problem of some other swarm intelligence algorithms that lead to suboptimal results [118]. In a performance comparison for unconstrained global optimization problems, WPA was found to possess superior performance in terms of accuracy, convergence speed, stability, and robustness when compared against PSO, GA, Artificial Bee Colony (ABC), ACO, and the Firefly Algorithm (FA), particularly for high-dimensional cost functions [119].

Artificial herding is another heuristic method particularly suited to human-swarm interactions. Inspired by shepherd's sheep-herding behavior, herding is advantageous for combining rule-based systems with learning-based algorithms such as deep reinforcement learning, which provides context-adaptation but suffers from scalability issue of the swarm [120]. Instead of training the

agents in the swarm with the control behavior, only the shepherd "agent" is trained to control part or all the rule-based swarm agents. In Nguyen et al. [120], the control of the swarm was done by inverse reinforcement learning, which allows for apprenticeship of the reward function from the expert training data that are human-generated policies.

An optimization algorithm closely related to the ACO with a faster convergence property is the ABC. Saied et al. [121] applied ABC for each UAV of the fleet to solve the distributed optimization-based control and collision avoidance of the UAVs.

The Fruit-Fly algorithm is a recent swarm intelligence optimization algorithm that models the food search behavior of fruit-fly swarm. Li et al. [122] used FA to solve the path-planning problem for UAV in a 3-D environment having complex terrain and multiple threat sites. The Fruit-Fly algorithm finds the optimum path that minimizes the cost function subject to length, threat probability, altitude, turning angle, climbing/gliding angle, and terrain constraints. This algorithm showed better performance compared to GA. Iscan et al. [123] proposed novel improvement of the classical Fruit-Fly algorithm to address the local minima convergence (premature convergence) by exploring not only the best solutions but the worst solutions in the search. Unlike other swarm intelligence algorithms, the Fruit-Fly algorithm is simple and straightforward to implement [122].

Spider-Monkey Optimization (SMO) is a state-of-the-art swarm intelligence optimization algorithm which models the fission-fusion social behavior of spider monkeys during food foraging. A study was conducted to compare SMO to other swarm intelligence algorithms for UAS path-planning problems [124]. Other swarm intelligence algorithms compared to SMO were ABC, Bat Algorithm, Grey Optimization, Harmony Search, FA, DE, Moth Search Algorithm, PSO, Whale Optimization Algorithm, and Greedy Crossover

Monarch Butterfly Optimization. The comparative study suggested that SMO provides the best path planning, has strong robustness, yields the most stable paths over multiple runs, and has the fastest convergence rates.

Chicken Swarm Optimization (CSO) is a state-of-the-art intelligent stochastic search algorithm inspired by chicken swarm behavior and hierarchical structure [125]. Key properties of CSO algorithm are good convergence speed and accuracy and good global optimization performance. Although the algorithm is susceptible to premature convergence, numerous modifications have been proposed to achieve global search ability. Liang et al. [126] proposed an Improved CSO for robot planning that has significant search efficiency and a high convergence rate. The simulation results were compared against the CSO and PSO. CSO has also been used for the parameter estimation of nonlinear systems [125].

We highlighted select state-of-the-art behavior and path-planning algorithms as well as fundamental algorithms for single and multiagents systems used in state-of-the-art robotics and autonomous systems.

3.5 INTELLIGENT STRATEGY

The hallmark capability of intelligent systems is reasoning for decision making. For the past decade, AI research focused on building human intelligence models mostly using reinforcement learning frameworks in complex environments, such as in strategy board games. Board and video games are a canonical testbed for AI algorithms and are opportune for machine intelligence development because they present conditions where the space of possible actions or moves to decide between is astronomically large, but the rules of the game are well-established, and production of large quantities of data is possible. Video games add complexity compared to board games by introducing partial observation of states (hidden states). Therefore,

choosing the optimal bet or strategy given the rules and available information is closer to human decision making. This section surveys significant applications of intelligent strategy and planning, namely in gaming scenarios.

3.5.1.1 Single-Agent Systems

Deepmind's AlphaGo

AlphaGo is a reinforcement learning based algorithm that plays the most complex strategy board game, the ancient Chinese game "Go." The game tree complexity of a 19 x 19 Go board is 10360 [127] compared to 10123 for the 8 x 8 chess board [128]. AlphaGo uses CNNs for perception and RL for decision making. The extensive training is divided into the SL from a database of 30 million expert moves known as the KGS dataset (maps the state to actions) and the self-play training.

The algorithm uses a Monte Carlo tree search (MCTS) governed by policy and value networks [129]. The MCTS is a heuristic look-ahead search method for a class of decision processes, such as in game trees. Instead of brute force expansion of the search tree to explore the best action trajectory or moves, it randomly selects and explores regions by simulated playout or rollout to determine the best moves iteratively using the reward operation of reinforcement learning. A balance between exploration and exploitation is key to choosing the best move from the current state in the search tree.

The significance of AlphaGo resides in the fact that it is the first computer program to win against a top professional human in the strategy game of Go, a feat in the category of high-complexity games, highlighting the potential to reach a superhuman level of learning with such an algorithm.

Deepmind's AlphaGo Zero

AlphaGo Zero, a notable departure from its predecessor AlphaGo, was modified to use the Residual Neural Network to predict both the

policy and value functions from a given state. Using only the raw board as input and rules, it was trained solely by reinforcement learning of self-play. Instead of MCTS rollouts, the Residual Neural Network is used to evaluate the best move.

The algorithm performs the following three main steps executed in parallel [130]:

1. In the self-play stage, a training set is created by selecting the best player version from 25,000 self-play games. At each move, the game state, the search probabilities from the MCTS, and the winner in the self-play are stored.
2. In the second stage, the network is retrained to optimize the weights from a minibatch of 2,048 positions sampled from the last 500,000 games. The Policy Loss function is a Cross-Entropy loss function between the network's predicted probabilities and the actual probabilities calculated from MCTS. The Value Loss is the mean-squared-error between the network's predicted value and the value computed from MCTS.
3. The third stage evaluates the latest network against the current best network. The evaluation selects the network with at least 55% wins out of the 400 games played between the two. Both agents represented by these two networks use MCTS to select their moves.

The significance of AlphaGo Zero is that it is possible to train an agent from zero capability (tabula rasa) to a superhuman level of expertise in complex settings using only a predefined set of rules. AlphaGo Zero had thrashed the previous version of AlphaGo in an impressive feat of mastery, with a score of 100 wins to none.

Deepmind's AlphaZero

AlphaZero is a generalized iteration of the latest AlphaGo Zero to two other strategy board games—chess and Shogi. The generalization removes the symmetry assumption associated with positions

in the Go board game and adds the feature of the possibility of the match to end in a draw. This iteration allows solutions to alternate turn-based games, transactions, or, more generally, scenarios in which:

- there are a fixed set of rules;
- there is a full observable state of agents, e.g., positions; and
- the primary goal of the opponent agent is known to prevent the other agent from winning.

Deepmind's MuZero

MuZero algorithm constitutes a leap forward in the capabilities of reinforcement learning. It does not rely on an available environment model, i.e., the game rules or game simulator. MuZero is the more generalized version of the iteration series of AlphaGo algorithms in that it learns an accurate model of an environment's dynamics, which it uses to plan the best possible actions [131]. Instead of learning from millions of degrees of freedom of interacting environment inputs, it models important aspects in the environment. This is important because the knowledge and simplicity of rules is not an option to a system dealing with the complexity of an actual real-world scenario.

The significance of MuZero for defense lies potentially in its application to simulated wargame scenarios with complex doctrine inputs and randomization. MuZero's ability to predict relevant quantities that allow for improved planning iteratively with no prior knowledge of the rules or environment dynamics to reach remarkably high performance is a state-of-the-art AI in strategy play and intelligent planning.

3.5.1.2 Multiagent Systems

OpenAI Five

OpenAI Five is a computer bot that plays a multiplayer battle arena video game called DOTA 2.

The underlying algorithm is Multiagent RL that uses self-play (180 years' worth of games against itself) to improve its capability.

In a live one-on-one battle game in 2017, the initial version played against a ranked professional player of the DOTA 2 and won. The last of subsequent versions was later demonstrated in a series of live five-on-five battles online over three days, ending with a win rate of 99.9% for the OpenAI.

The significance of OpenAI Five's performance in DOTA 2 demonstrates success in an imperfect information, zero-sum game environment that combines several attributes, such as competitive actions, collaborative actions, and dynamic environment characteristics, all of which find parallels in real-world conflict. Contrary to the environments in which the single agent systems like AlphaGo and its variants operate, the real-time requirements of fast decision making, intent prediction and anticipatory actions, continuous action and observation space, long-term play, and partially observed information are what make OpenAI Five a worthy AI candidate in multiagent competitive and collaborative settings.

Deepmind's AlphaStar

AlphaStar is a computer bot based on the Multiagent RL algorithm that plays the real-time strategy video game StarCraft II. This video game environment is one in which perfect information for planning is not available and fast decision making is required, as time-dependent events may change the current best strategy of a given time.

AlphaStar is trained initially from thousands of human replays (supervised learning). Using policy-gradient-based RL, the parameters are trained so the win rate is maximized. Self-imitation (UPGO) is used to deal with off-policy updates [132]. AlphaStar incorporates multiple feature extraction stages. Minimaps are extracted with an RNN. Time sequence of observations is processed by an LSTM. Scatter connections are used to incorporate spatial and nonspatial information.

The vast combinatorial action space is addressed using an autoregressive policy structure and a recurrent point network.

Similar to OpenAI Five, the significance of AlphaStar is the resultant strategy the bot employed over the course of the competition, which displayed several remarkable abilities from the experts' point of view based on their analysis and interpretations of the games. Among the observations are the initial aggressive and confrontational play due to prioritizing short-term rewards over long-term ones, the unconventional wisdom instigating a suboptimal play of the opponent, and opponent intent or vision prediction [132, 133].

3.6 OPPONENT MODELING AND WARGAMES

Opponent modeling is the prediction of the behavior or strategy of one or more agents having full or partial states in an adversarial game by using prior knowledge and observations. Opponent modeling is used to predict rationality-based actions of an unknown opponent and is therefore very important for defense-related wargaming modeling and simulation. Opponent modeling addresses the explicit interpretability of an opponent to model the opponent's behavior. This contrasts with systems like AlphaGo and its variants, where the internal representation of opponent strategies is inferred indirectly.

Opponent modeling is categorized by three approaches—strategy classification, goal-based generative models, and policy approximation [134]:

1. Strategy classifications aim to predict the strategy from learned experts via supervised or game-theoretic methods.
2. In goal-based generative models, how one or more agents achieve a goal is the basis for the behavior classification.
3. Policy approximation learns the function mapping from states/action set, which estimates the true policy behind the opponent's behavior.

While strategy classification methods are abundant in the recent literature and favored over goal-based generative model methods, they both suffer from several common drawbacks, such as the following:

- Reliance on the process of translating expert-designed semantic labels to action distribution over future game states.
- Potential disagreement between labeled opponent models and the actual opponent's internal representation affecting the performance of discriminative opponent models.
- Reliance of the discriminative models on expert knowledge.
- Scarcity of knowledge about the opponent affecting the performance of the models.

In contrast to discriminative and generative modeling, policy approximation provides the distinct advantage that it produces the probability of future game states. Policy approximation will be the focus herein concerning opponent modeling.

The policy approach can be approximated using NNs (represented as decision trees) or state-action tables. Expert domain knowledge can still be incorporated into the policy approximation framework in the form of state/action sets.

Tang et al. [135] presented a novel AI algorithm for a two-player fighting game based on opponent modeling combined with evolutionary strategy to determine the best response against the opponent. The evolutionary strategy used an Enhanced Rolling Horizon Evolution Algorithm (E-RHEA). The opponent modeling was performed via three variants—on-line supervised learning with cross-entropy loss function, policy-gradient reinforcement learning, and Q-learning-based reinforcement learning. The variant algorithms were investigated and showed significant improvement to RHEA-based strategy-only play. It is suggested that RHEA with opponent model outperformed state-of-the-art MCTS-based

fight bots. The E-RHEA based on policy-gradient reinforcement learning won first place at the 2019 Fighting Game AI Competition and 2020 IEEE Conference on Games.

Synnaeve et al. [136] combined CNN and RNN to conduct opponent modeling by exploiting spatio-sequential correlations in large training data from the StarCraft (Brood War) games. Real-time strategy was computed from the encoder-decoder RNN architecture to provide state estimation and future state predictions from only previous and partial observations of the game dynamics. The encoder used a concatenated input of learned embedding of both players' action and state features. Two encoders were examined—ConvNet and Convolutional LSTM. The embedding from the encoders passed through recurrent LSTM cells, which allowed the capture of information from previous frames. Analysis of off-line testing showed that the encoder-decoder, LSTM-based architecture performed better at predicting current and future states than rule-based baselines.

Similarly, Chen et al. [137] proposed a novel Bayesian Policy Reuse (BPR) approach for nonstationary opponents in Markov games. The algorithm formed a policy library in an off-line learning phase. Unlike the MinMax-Q-based algorithms that selected conservative actions, the selected actions of the suggested algorithm considered the opponent behavior model. Opponent modeling was done via the NN approximator, which updated opponent models with sampled observation. The network was trained with the opponent behavior state and observation tuples to provide the probabilities of the candidate actions as output. A belief function was used to measure the similarity between each approximated opponent policy. The algorithm combined BPR and opponent models to help the agents reuse the best response policy. The novel algorithm was tested in a soccer domain with two agents and showed better performance compared to existing approaches.

In nonstationary scenarios, such as in zero-sum Markov games where the opponent strategy changes concurrently, Chen et al. [138] proposed an opponent modeling stage and action heuristics integrated with an eXtended Classifier System (XCS) to enhance the action selection and policy learning. The opponent model was constructed to capture the game interactions and propagate the predicted opponent behavior for action selection and classifier evolution. To speed up the learning process, an accuracy-based eligibility trace mechanism was employed via reinforcing the XCS matching of historical traces according to their accuracy. By considering the accuracy of each classifier in the action heuristic function, the algorithm avoided the associated limitations with assuming that the opponent followed optimal strategy, such as in MinMax based Q-learning algorithms. The simulated Markov games demonstrated efficient and reasonable use of the heuristic policies for single-agent heuristics. For multiheuristic, multiagent games subject to Pareto optimal action selection strategy, the proposed algorithm required more time to update the classifiers.

Besides concurrent opponent strategy changes, nonstationary scenarios can also manifest in deceptive measures by deliberate, inconsistent dynamics, conveying false cues or concealing true information. An important contribution in a competitive nonstationary, multirobot environment, Chen and Arkin [139] proposed a novel countermisdirection approach for behavior-based, multirobot teams. They introduced a countermisdirection agent (CMA) to detect the misdirection process and act collaboratively to stop it.

The robots' groups in the scenario were the Mark group, misdirection team, and countermisdirection team. The mark group is the subject of misdirection and modeled on each of the mark agents in the group using the Granovetter threshold model. This model captures the behavior of the flock of animals

that can be misled or stopped. The misdirection team seeks to misdirect the mark group. The misdirection team is characterized by the leader of the group, who triggers the misdirection process by initiating action from the Wander Behavior model; the rest of the group follows to form a group behavior that misdirects the mark group. The CMA group has no prior knowledge of the misdirecting team's leader goal location. The CMAs cannot identify marks from misdirection team agents. The CMA monitors the movement of the active agents based on its threshold and within its observation range for a period of time. They estimate an intercept position in front of the marks in motion based on the collective movement of the active agents. The CMAs then move to the estimated intercept position and form a barrier to prevent the marks from reaching the misdirecting team leader position. This algorithm simulated results that have potential in the field of mobile robotics with military applications.

In the multiagent domain, Zheng et al. [140] proposed an Extended Deep Bayesian Policy Reuse (Deep BPR+) algorithm to handle nonstationary opponents in Markov games. Opponent policy was simultaneously determined using the belief model based on rewards signals and opponent models. The Opponent model was obtained from a DNN approximator instead of tabular representation using past sequence of moves. It was obtained by maximizing the log probability of the approximated opponent policy with an added regularizing entropy term of the policy to address the overfitting issue. To measure the similarity between the opponent's different policies, Kullback-Leibler divergence was used. To speed the online learning process, the Distilled Policy Network (DPN) was used to initialize a starting policy and combine multiple response policies into one, therefore improving response time. DPN offered a generalized way of initializing the starting policy, leading to better performance without being concerned about the choice of response policy.

He et al. [141] presented an implicit multiagent opponent modeling with optionally enabled explicit modeling, based on the deep Q-Network (DQN), using Mixture-of-Expert architecture. The suggested model demonstrated automatic learning of different strategy patterns of the opponent in a simulated soccer game and popular trivia game.

Biro and Walker [142] applied the RL approach to play calling in football. The algorithm finds the utilities and optimal policy from value iteration and greedy policy computation at each state, which provides valuable insights and allows for an improved comprehension of the game.

Opponent modeling is an area of active research, as there are several challenges that pertain to multiagent settings—particularly, how an action taken by one agent at a given time jointly changes the environment state and the outcome of actions taken by other agents.

3.7 COGNITIVE ELECTRONIC WARFARE

The electromagnetic spectrum is an important front of warfare; modern and conventional weapons rely on electromagnetic, electro-optical, and acoustic signals for communication, intelligence, sensing, and weapon delivery. The emergence of modern AI methods has brought about opportunities to reframe traditional electronic warfare (EW) problems in statistical signal processing and information theory as AI problems. The use of AI in EW, now termed cognitive EW (CEW), has been gaining immense attention, particularly in areas of automatic modulation classification, automatic intrapulse modulation classification, and radar pulse repetition interval tracking [143, 144].

Typical modern EW systems fall under three fields of applications or system component processes [140] (see Figure 3-6):

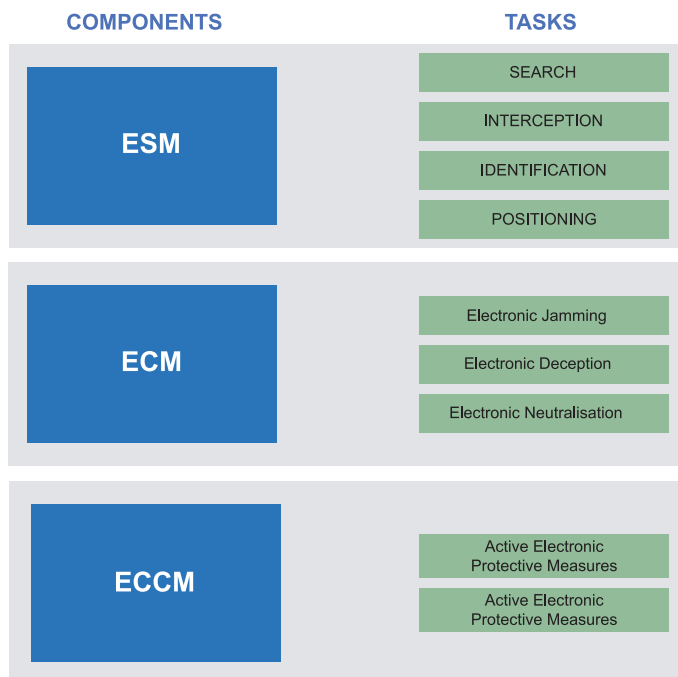


Figure 3-6. EW System Components (Source: QinetiQ).

1. EM threat interception, recognition, and location are provided by an electronic support measures (ESM) system.
2. Offensive measures provided by an electronic countermeasure (ECM) system disturb or deny an enemy’s EM operational means.
3. Defensive measures provided by an electronic counter-countermeasures (ECCM) system provide resiliency and protection against the enemy’s ECM.

3.7.1 Electronic Support Measures

AI techniques using residual networks from deep-learning algorithms have been successfully used for ESM to perform automatic classification of signals. Ammar et al. [143] have shown the ability to classify signals of 24 types of modulations.

Liao et al. [145] devised an unsupervised scheme for classifying and identifying received radar pulse signals using the Ward Clustering method and Probabilistic Neural Network (PNN). The scheme used adaptive filtering and spectrum

analysis in the data processing stage, followed by a preclassification state in which Ward's clustering was employed, the output of which was fed to the accurate classification stage. The latter used the preclassified signals for training the PNN. The Calinski-Harabasz index, Gap index, Silhouette index, and Davies Bouldin index were used as validation criteria for determining the optimal number of clusters. The PNN was further optimized by the maximum of the classification validity index. To increase the identification accuracy of the PNN classifier, the bivariate correlation analysis was used. The scheme showed better classification stability and achieved 100% for classification and identification.

In automatic intrapulse modulation classification, Qu et al. [146] used time-frequency analysis, image processing, and CNN to recognize 12 radar signals. To extract time-frequency images for CNN, Cohen Class Time-Frequency Distribution (CTFD) was used with the Gaussian function as a kernel, followed by a series of image-processing steps to remove the background noise. The CTFD was advantageous over Short Time Fourier Transform in terms of adaptability and over Wigner-Ville Distribution in terms of high time-frequency resolution and removal of cross-terms that appear in nonlinear frequency modulations at higher time-frequency resolutions. The recognition stage used LeNet-5 based CNN with the RELU activation function. The confusion matrix results of the recognition stage showed a probability of successful recognition of more than 96.1% for 12 different radar signal modulations of -6 dB SNR.

Supervised techniques from computer vision for visual detection of signals use the spectrogram waterfall images to divulge signal features and characteristics. O'Shea et al. [147] demonstrated signal detection and localization using a variation of YOLO's (Tiny-Yolo) region-based CNN that was trained on a 20k event containing spectrograms.

Lee et al. [148] compared forward DNN and LSTM structures for predicting radar-jamming signals. They showed that forward DNN required the extraction of features from pulse description word list of radar signals prior to training, while LSTM structures directly used the list for training and prediction. The LSTM showed greater prediction accuracy on average than the forward DNN structure, with the caveat that the training took longer for LSTM. It was suggested that forward DNN and LSTM could be used effectively to predict unknown radar signals with an average accuracy of 92% or higher.

A recent example of cognitive EW systems in development is BAE system's Controllable Hardware Integration for Machine-learning Enabled Real-time Adaptivity (CHIMERA). BAE is developing a reconfigurable AI-powered hardware platform that enables advanced deciphering and detection capabilities against electronic assaults. The CHIMERA is expected to adapt its radio frequency (RF) configuration and signal features in real-time [149].

3.7.2 ECMs

Qiang et al. [150] proposed a jamming-style selection algorithm based on support vector machines using a jamming-rule base. The jamming-rule base is constructed from the features of airborne multifunctional radars with cognitive capabilities in air-air combat. The feature space used is Pulse Repetition Frequency, Carrier Frequency, Pulse Width, Coefficient of Resemblance, and Box Dimension. The proposed method does not depend on the radar state estimation and is advantageous in terms of fast convergence and real-time requirement compared to online learning methods like Q-learning.

One of the techniques used in ECM is the deceptive transmission of pulses at the pulse interval of enemy radar signal. The pulse tracking of a pulse repetition interval (PRI) requires the continuous

prediction of time of arrival of pulse trains. In ESM, PRI tracking extracts specific pulses from interleaved pulses. In ECCM, a PRI predictor is used to discern signal pattern and counter decoy sources [143].

Autoregressive RNN like LSTM have been used recently in generating a radar pulses stream that belongs to the same pulse probability distribution of the received pulse stream. The ability to generate similar pulses may be applied to the interpulse extrapolation for tracking as part of system functionality in ESM. This may also be applied in ECM, where LSTM could be also used for saturating the spectrum with fake returns to raise enemy's radar probability of false alarm by synthesizing train of pulse sequences in time [143].

Mission-critical applications require fast adaptation to cognitive jammers and dynamic environments. Amuru et al. [151] proposed a Multiarmed Bandit-Based learning algorithm that optimally jams malicious transmitter-receiver pairs without an a priori knowledge about its strategy or channel gains. The algorithm was tested against static and adaptive strategies and showed the capacity to track different strategies used by adaptive transmitter-receiver pairs.

The U.S. Air Force plans to take intelligence warfare to the electronic spectrum via its Kaiju project. The project seeks to create an EW ecosystem to demonstrate next generation of systems capable of autonomous electronic attacks and enemy air defenses countermeasures.

3.7.3 ECCMs

In ECCM, cognitive radio (CR) is an automatic spectrum allocation system that locates appropriate channels in the spectrum based on the perceived conditions in the environment, such as spectrum interferences and jamming. Cognitive radars employing CRs have a high countermeasure performance against interferences and jamming [143]. The adaptability of CR to changes in the

environment presents challenges to the EW system to intercept, track, and jam its signals. A countermeasure is the use of an RL-based adaptive system.

Yangyang et al. [152] developed an intelligent jamming method based on RL and analyzed the performance of a DRL-based, antijamming algorithm as a measure against the intelligent jammer in various communication and jamming modes. Simulation results suggested that the RL-based jammer could effectively restrict the performance of the DRL-based, antijamming method.

Kang and Bo [153] proposed a DQN method for the frequency-hopping strategy of a cognitive radar without knowledge of a jamming model. This reinforcement-learning-based method allows the cognitive radar to learn the jammer's strategies based on the changes it perceives and then adopts an optimal strategy of avoiding jamming signals.

SECTION 04

SYSTEMS AND PROGRAMS APPLYING AI FOR WEAPONS SYSTEMS

In this section, we briefly cover examples of the various ongoing systems and programs relating to AI for weapons systems. We mainly focus on U.S.-based systems and programs but conclude with two subsections in Section 6 highlighting related developments by Russia and China.

The U.S. government has established a few important institutions specifically focused on AI. Within the DoD is the Joint Artificial Intelligence Center (JAIC) and the Under Secretary of Defense for Research and Engineering (USD[R&E]) community of interest (Col) on autonomy.

The JAIC was established by the DoD in 2018 to expand and transform U.S. military capabilities and services by leveraging AI technology. It aims to accelerate the adoption of AI as well as lead the national effort to maintain the U.S. military technological edge and as a global force in the next-generation technologies related to AI. According to a memo released in December 2021, the JAIC will be merged under the Chief Digital and Artificial Intelligence Officer together with the Chief Data Officer and Defense Digital Service [4].

The USD[R&E] stood up the Cols, one of which focuses on autonomy. The Col's are designed to encourage multiagency coordination and collaboration in cross-cutting technology focus areas with broad multiple-component investment. Cols provide a forum for coordinating scientific and technology strategies across the Department, sharing new ideas, technical directions, and technology opportunities; jointly planning

programs; measuring technical progress; and reporting on the general state of health for specific technology areas. The autonomy Col has four subgroups focused on the following:

1. **Human/Autonomous System Interaction and Collaboration:** The keys to maximizing the human-agent interaction are instilling confidence and trust among the team members; understanding each member's tasks, intentions, capabilities, and progress; and ensuring effective and timely communication. All must be provided within a flexible architecture for autonomy, facilitating different levels of authority, control, and collaboration.
2. **Machine Perception, Reasoning, and Intelligence:** Perception, reasoning, and intelligence allows for entities to have existence, intent, relationships, and understanding in the battlespace relative to a mission.
3. **Scalable Teaming of Autonomous Systems:** Collaborative teaming is a fundamental paradigm shift for future autonomous systems. Such teams are envisioned to be heterogeneous in size, mobility, power, and capability.
4. **Test, Evaluation, Validation, and Verification:** The creation of developmental and operational T&E techniques that focus on the unique challenges of autonomy, including state-space explosion, unpredictable environments, emergent behavior, and human-machine communication.

On the U.S. federal level and outside the DoD are important AI-related groups, such as the National Artificial Intelligence Advisory Committee (NAIAC), the National Security Commission on Artificial Intelligence (NSCAI), and the National Artificial Intelligence Research Institutes.

The NAIAC is tasked with provisioning critical information, recommendations, and reports to the President and AI Initiative Office on topics related to the national AI initiative. The NSCAI was established in 2018 as an independent commission by Congress to issue reports to the President and Congress on the advancement of AI, related ML developments, and associated technologies. The commission reviews U.S. AI's technology advances and competitiveness related to national security, defense, and investments and makes recommendations in developments, education, ethics, standards, and technology management in relation to AI.

National Artificial Intelligence Research Institutes is a program born out of National Artificial Intelligence Research and Development Strategic Plan whose purpose is to enable long-term research and U.S. leadership in AI through the creation of joint government and industry-funded AI research institutes. This program is a joint government effort between the National Science Foundation, U.S. Department of Agriculture's National Institute of Food and Agriculture, U.S. Department of Education's Institute of Education Sciences, U.S. Department of Homeland Security's Science & Technology Directorate, NIST, DoD's Office of the Under Secretary of Defense for Research and Engineering, and IBM Corporation.

In the next sections, we focus on specific programs and systems in development at the intersection of AI, autonomous systems, and weapons systems. Our survey of related programs and systems is far from comprehensive but is more of an overview of those AI-enabled systems (or related programs) that have battlefield applicability to weapons systems. We have grouped these systems by their domain (aerial, maritime, and land), collated those

AI-systems designed for swarming purposes, and listed AI-enabled tools for battle management, navigation, and targeting.

4.1 AERIAL SYSTEMS

4.1.1 Next-Generation Air Dominance (NGAD) Programs

The U.S. Air Force's development of sixth-generation fighters is currently being developed under the NGAD program. One key element of these fighters is the integration of the "AI wingman," the AI-enabled autonomous functions and manned-unmanned teaming (MUM-T) capabilities.

4.1.2 Shield AI Hivemind (Heron Systems)

Heron Systems' AI is the victor of the DARPA's AlphaDogfight competition. Its impressive show of intelligent maneuverability, combat dynamics, and intent prediction resulted in a win against multiple, highly capable AIs in the competition and against a U.S. Air Force F-16 fighter pilot in a simulated dogfight. The remarkable win echoed throughout the DoD, U.S. Armed Forces, defense, intelligence, and academic communities. According to Heron systems' Brett Darcey, the dominant performance of Heron Systems' Falco agent (AI) will be further set apart by its potential use of a Shield AI (AI and robotics company) platform capable of operating in GPS and communication-denied environments [154]. The system is being integrated and planned to be scaled, operationalized, and fielded into existing unmanned platforms and in NGAD's unmanned combat aerial systems (UCASs).

4.1.3 Shield AI V-Bat

The V-Bat is a versatile vertical takeoff landing, fixed-wing autonomous vehicle (Figure 4-1). The V-Bat can carry out a wide range of mission operations, such as infantry clearance, air-defense breach, and swarm operations, thanks to its integrated Shield AI's Hivemind autonomy core [155]. The Hivemind autonomy core negates the need for human operator, GPS, and RF links.



Figure 4-1. V-Bat UAV in Transition Flight From Vertical Takeoff
(Source: Shield AI Used With Permission, <https://shield.ai/products>).

4.1.4 Kratos XQ-58 Valkyrie

The U.S. Air Force's Skyborg program provides rapid prototyping via an AI-enabled, open-autonomy architecture that emphasizes continued scalability, modularity, commonality, and portability. It also seeks to develop low-cost attritable aircraft technology-based UCAS (Figure 4-2). As one of the U.S. Air Force's Vanguard programs, Skyborg will fast-track the capabilities toward MUM-T and resilient autonomy, thereby improving survivability and lethality against near-peer adversaries [156].



Figure 4-2. U.S. Air Force's Skyborg Conceptual Design (Source: AFRL, <https://www.af.mil/News/Article-Display/Article/1796930/skyborg-program-seeks-industry-input-for-artificial-intelligence-initiative/>).

The U.S. Air Force has conducted several flight tests of Skyborg AI aboard Kratos' XQ-58 Valkyrie to demonstrate the Skyborg open and modular autonomy architecture. At the latest test on March 2021, the demonstrator released the Altius-600 UAS payload [157].

4.1.5 MQ-20 Avenger UCAS

The Skyborg team demonstrated the portability and scalability of the Skyborg autonomy AI on a different platform, the MQ-20 Avenger UCAS [158]. The modularity of the Skyborg autonomy AI core allows for autonomous capability increments and more advanced AI integration.

4.1.6 Autonomous Loitering Munitions

Autonomous loitering munitions are a category of aerial weapons. They are launched, and then the munition weapons system waits until its required use to strike. These are often lightweight, multidomain capable, and ground launched. Recent work has been done to add the capability for these weapons systems to be fully autonomous. Examples of integrating autonomy into loitering munitions include IAI's Harpy, KARGU, ASN-301, and Orbiter 1K. Each have a "fire and forget" feature where they can be launched and then acquire and prosecute targets without further supervision.

4.1.7 Dynetics X-61 Gremlins

X-61 Gremlins is a technology demonstration platform originating from DARPA's Gremlins program designed to provide a semiautonomous, mid-air recoverable, and low-cost UCAS (Figure 4-3). The X-61 can carry a variety of payloads, including EO/IR imaging, EW sensors, and weapons.

4.2 MARITIME SYSTEMS

The U.S. Navy is applying AI technologies to several ongoing autonomous platform development programs. Concordant to the DoD directive guidelines, AI technology is employed through



Figure 4-3. Dynetics X-61 Gremlins UCAS (Source: DARPA, <https://www.darpa.mil/news-events/2020-01-17>).

the Navy's Unmanned Maritime Autonomy Architecture and Common Control System.

The Sea Hunter is an unmanned Navy platform operated to develop TTPs. It serves as one of the Navy's test platforms for autonomy development. Its autonomy is so sufficiently developed that it navigated from San Diego to Pearl Harbor. A second Sea Hunter, called Seahawk, was launched in August 2020 (Figure 4-4).



Figure 4-4. U.S. Navy's Sea Hunter in Its First Demonstration (Source: U.S. Navy, https://www.navy.mil/Portals/1/Strategic/20210315%20Unmanned%20Campaign_Final_LowRes.pdf?ver=LtCZ-BPIWki6vCBTdgtDMA%3D%3D).

4.3 LAND SYSTEMS

4.3.1 QinetiQ/Pratt Miller's Expeditionary Autonomous Modular Vehicle (EMAV)

Jointly developed with QinetiQ North America (QNA), EMAV is a large robotic combat vehicle

(RCV-L) (Figure 4-5) designed to meet the U.S. Army RCV program's specific requirements for the RCV decisive lethality continuum concept. The EMAV houses QNA's modular ground autonomy core open architecture. The system is being verified for its support capabilities in a MUM-T operational environment.



Figure 4-5. QinetiQ/Pratt Miller's RCV-L (Source: QinetiQ [159]).

4.3.2 Textron Systems' Ripsaw M5

The Ripsaw M5 is a fifth-generation medium RCV (RCV-M). The electric tank, a U.S. Army technology demonstrator with fully autonomous capabilities, is part of its three-tier RCV decisive lethality continuum concept (Figure 4-6).

4.3.3 Rheinmetall's Lynx KF41

The Lynx KF41 is an optionally piloted infantry fighting vehicle under development that integrates a large suite of Raytheon's sensors and AI-enabled autonomous capability. The infantry fighting vehicle will be equipped with a virtual crew AI that provides continued scanning and detection of battlefield landscape for increased situational awareness and automatic target recognition (ATR) to alert the crew for a swift tactical course of action.

4.4 SWARM SYSTEMS

Swarm systems are fully (operator is still on the loop) or partially autonomous systems capable of coordinating autonomously to execute swarm-



Figure 4-6. Army's Ripsaw M5 Unmanned Battle Vehicle
(Source: U.S. Army Photo Courtesy of Textron Systems, <https://asc.army.mil/web/news-rapid-robotic-requirement-relay/>).

based missions (e.g., flocking, rendezvous, and “play calling”). These are systems that operate together to achieve a unified end. The DoD is actively pursuing the development of such systems.

4.4.1 DARPA's Offensive Swarm-Enabled Tactics (OFFSET)

DARPA's OFFSET is a program enabler of the next-generation ecosystem of combat technologies that seeks to produce a diverse set of swarm mission-capable unmanned autonomous vehicles. The program envisions using human-swarm teaming of upwards of 250 unmanned aerial and ground systems.

4.4.2 Collaborative Small-Diameter Bomb Swarm (CSDB)

CSDB is a swarm technology based on autonomous weapons munitions in development under the Golden Horde program, which falls under the umbrella of the U.S. Air Force's Vanguard program whose objective is to expedite the advancement

of warfare proven and effective AI technologies. The first phase of the Golden Horde program culminated in the development and demonstration of a swarm of CSDBs operating collaboratively to engage a target [160]. The subsequent phases build upon the demonstration in a digital weapons ecosystem, where the virtual enhancements are carried and evaluated in a virtual simulation environment dubbed “Golden Horse Colosseum.”

The swarm technology of the CSDB seeks to provide autonomous identification, target selection, and strike capabilities. The autonomous behavior of CSDBs modeled after the collaborative autonomy of play calling in which the collaborative play is chosen from a preestablished set of plays. The decentralized control means that if one or more group members are lost, the mission can be still completed since each group's members runs the autonomy core [161].

4.4.3 Perdix Swarm

Perdix is a technology demonstrator of an autonomous multiagent weapons system of microdrones with decentralized control that operate swarm reconnaissance and other tactical missions. The decentralized control, dubbed “hive-mind,” is a self-adaptive and expandable system with no centralized or human control. Perdix technology shows itself to serve as an anti-air-defense weapon. Each microdrone in the swarm carries a transmitter for jamming seen by enemy air defenses as a decoy that, in large numbers, overwhelms the radar systems of the enemy air defenses [162].

4.4.4 Mako UTAP22

Developed by Kratos Unmanned Systems Division, the Mako UTAP22 platform is being developed to specifically operate in tactical MUM-T and carry out swarm operations. The UCAS embeds the U.S. Air Force Skyborg's autonomy core and can coordinate commands for attack and maneuvering from ground-based and air-based command and control.

The UCAS's payloads include internal and external weapons payloads and advanced EW and sensor wingtip pods [163].

4.4.5 Coyote UAS Block 3

Coyote is a tube-launched, small autonomous, expandable, and multimission UAS adaptable for diverse individual or swarm missions, including EW, surveillance, and strike [164]. Developed by Raytheon Missile & Defense, the UAS can be launched from the ground, air, and sea. The U.S. Army has already successfully tested Coyote UAS Block 3 against a UAS swarm, but due to the urgent need for more performant counter-UAS solutions, variants of the Coyote are being further enhanced to fly faster and farther. The U.S. Navy seeks to field its own version launched from unmanned surface vehicles (USVs) and UUVs.

4.4.6 Control Architecture for Robotic Agent Command and Sensing (CARACaS) Swarm

CARACaS is an autonomy core framework for multiagent coordination and control that has been used by the Navy to create "swarmboats" (Figure 4-7). The CARACaS system provides fully distributed, multiagent operations. The cooperative behavior models supported by the autonomy core are patrol, track, inspect, and trail. The system is planned for heterogenous platform teaming and adaptable to attrition and subsystem failures [165, 166].

4.4.7 Riptide Swarm Micro-UUVs

The Riptide family of UUVs is a series of highly flexible robotic underwater platforms capable of long-range, high-endurance, and high-speed critical missions at a great depth. Draper Labs, in collaboration with Riptide Autonomous Solutions and Teledyne Benthos, is developing a swarm system of micro-UUVs for the Navy [167] (Figure 4-8). Thanks to the integrated maritime open architecture autonomy that Draper developed for the U.S. government's open architecture for mission autonomy, the swarm agent can carry out autonomous sorties, maintain underwater communications, surface for periodic GPS updates, and exchange acoustic target intelligence data to engage surface targets.

4.5 BATTLE MANAGEMENT AND INTELLIGENT COMMAND & CONTROL

The DoD continues its efforts in broadening the capabilities of the U.S. Armed Forces via the exploration of effective symbiotic human and machine operations through MUM-T, a mission autonomy paradigm viewed as a leap departure from the conventional warfare operations. This paradigm synergizes human intuition and oversight, the effectiveness of autonomous systems with adaptation, resilience, and on-demand decision-guiding abilities of AI to deliver highly accurate mission execution, operational flexibility, and increased lethality and survivability. MUM-T



Figure 4-7. Swarm II USVs With CARACaS During Demonstrations (Source: NASA [165]).



Figure 4-8. Riptide Autonomous Solutions Micro-UUVs
(Source: Navy, https://www.navsea.navy.mil/ortals/103/Documents/NUWC_Newport/ANTXdocs/ANTX_Playbill_LR_15AUG.pdf?ver=2018-08-22-105459-353).

constitutes a shift toward the beginning of an era of human-machine cooperation.

To meet the data-intensive and network-processing requirements of the next-generation battlefield, the DoD is ramping up development of a plurality of systems supporting the Joint All Domain Command and Control with AI-enabled mission planning and combat management capabilities.

Raytheon is developing a battlefield decision support tool dubbed “ARAKNID” (Anytime Reasoning and Analysis for Kill-Web Negotiation and Instantiation across Domains) as part of the DARPA Adapting Cross-Domain Kill-Webs program. It was recently used in a multinational exercise led by U.S. Northern Command called Global Information Dominance Experiment 3, or GIDE 3, where ARAKNID took in a variety of information and then recommended operationally relevant courses of action to human decision makers.

BAE continues to develop a semiautonomous mission-planning tool under the names of Multi-domain Adaptive Request Service (MARS) and, more recently, the Adapting Cross-Domain Kill-Web (ACK) program. MARS was a Phase 1 effort completed in partnership with Carnegie Mellon University and Uncharted Software. Together they developed the MARS software, which allows quick mission plan updates given new target information. ACK is the continuation of this work under a Phase 2 effort with DARPA, with the goal of demonstrating this technology in a multidomain, operationally relevant environment.

The U.S. Army’s Constructive Machine-learning Battle with Adversary Tactics (COMBAT) program seeks to use AI to develop advanced adversarial war game tactics to stimulate U.S. Forces countermeasures and retaliatory tactics in a way that will ensure long-term tactical advantage of the U.S. Forces and enforce unforeseen battlefield reality for the enemy.

DARPA’s Explainable AI program will aid human, battle management decision makers by ensuring that the AI can produce explainable models and reasoning and present the rationale in ways that a human can understand. This process will use neurosymbolic AI and state-of-the-art human-computer interface techniques.

4.6 ISR & TARGETING SYSTEMS

4.6.1 SRC’s HPEC Pod

The High-Performance Embedded Computing (HPEC) pod is an AI-enhanced ISR payload designated for the Agile Condor (MQ-9 Reaper variant, Figure 4-9) addressing the requirement for processing, exploitation, and dissemination (PED) of data for the next-generation, semiautonomous unmanned combat aerial vehicles (UCAVs). The HPEC pod hosts a modular AI suite capable of near real-time autonomous processing of target acquisition, enhanced situational awareness, and dissemination of valuable combat information.



Figure 4-9. MQ-9 Reaper in Flight (Source: Photo by Airman 1st Class William Rosado, <https://www.dvidshub.net/image/5625956/mq-9-reaper-flight>).

SRC Inc.'s AI suite boasts adaptive decision making and provides multimode mission continuity in communications-denied environments.

4.6.2 Nemesis

While augmented reality provides the information advantage and enhances the perception of the Warfighter during combat, HRL Labs' NEuroMorphic

EyeS In the Sky (NEMESIS) adds the tactical dimension of automated, critical decision making. NEMESIS is a bioinspired visual AI system (Figure 4-10) that seeks to emulate the human vision and cognition by fusing a multimodal sensor, recognizing scenes and situations, interpreting context, and proving fast, intelligent decision and tactics in real-time to the combatant. In a series of two demonstrations, NEMESIS first tested its activity recognition and situational awareness capabilities. In the second test, mission analytics were demonstrated, such as mission phases, incorrect actions, and mission anomalies [168].

The Maven Project, initiated by Office of the Secretary of Defense, seeks to catalyze the next generation AI/ML technologies that will automate PED for tactical UAS and Mid-Altitude and High-Altitude ISR platforms by exploiting the large volume of field data and Full Motion Video (FVM) and generating actionable intelligence at a faster pace. The initiative is to develop algorithms for ATR using deep learning and computer vision for detection/identification/tracking and AI for



Figure 4-10. HRL's Nemesis Eye Interprets the Battlefield for Actionable Intelligence (Source: Getty Images, <https://www.gettyimages.ca/detail/photo/new-technologies-to-be-used-in-specialized-military-royalty-free-image/1010598692>).

ascertaining the proper course of actions for further stages of decision making.

4.7 NAVIGATION

Conventional weapons systems rely on a GPS signal locked inertial navigation system solution to get accurate real-time position, navigation, and timing (PNT) information. However, in austere and electronically hostile environments where communication signals might be subjected to jamming and spoofing, mission execution might be disrupted due to a loss of positional and timing accuracy in the navigation systems. AI-driven technologies are being developed and tested to maintain navigation in such environments.

Northrop Grumman is developing new PNT technology using AI and ML algorithms to detect and identify GPS threats in low-power RF signals near the noise floor that will affect navigation by searching large RF spectrum data for threat characteristics and anomalies. The algorithm of the GPS threat detector automatically updates its own list of reference threats and communicates the threat information to other combat systems or operators within the network. The detection allows for the mitigation of GPS degradation by reconfiguring its operation [169].

In addition to the invisible electronic threats, the potential use of long-range missiles to target satellite constellations has become a reality since 2007, when China successfully shot down an orbiting satellite [170]. In view of such threats, the government contracted with KBR, Inc., to research and prototype an alternative PNT navigation system for use in GPS-denied environments. The Stealth and Cognitive Agile Navigation System project will produce a cognitive PNT navigation system that will provide continuous real-time positioning using AI and cognitive dead-reckoning models with enhanced environment sensing [171].

Navigation in a deep ocean environment remains an unexplored and unexploited area as a warfare

entry due to the technical challenges associated with loss of surface and satellite communications when underwater. Therefore, a growing area of research is exploring the use of AI algorithms to overcome GPS-deprived navigation as in the deep sea. DARPA's Angler program aims to develop an underwater platform capable of a fully autonomous, long-duration mission deprived of satellite and surface communication with the help of a sensor suite providing perception in dark, turbulent, and semiopaque surroundings.

This Page Intentionally Left Blank

SECTION 05

AI IN FUTURE COMBAT

U.S. defense strategists are betting on AI-enabled weapons and sensors to project complexities onto the battlefield that will overrun and overwhelm adversary systems, defenses, and combat decision making [172]. Mosaic warfare is the terminology used to capture this strategic vision for future combat. Former director of DARPA's Strategic Technology Office stated that the goal:

"... is to create the interfaces, communications links, and the precision navigation and timing software – the technology backbone – to allow these exquisite systems to work together. On PowerPoint illustrations of battlefields,

these communication links are often portrayed with lightning bolts. One of our mottos is to make lightning bolts real." [173].

Mosaic warfare concept is a warfighting strategy that employs multidomain, nonuniform, and asymmetric force projection that employs a variety of weapons platforms of different types, classes, configurations, and sizes. By the mosaic tile analogy, the concept foresees all the constituents coming together to create an unconventional strike force that disrupts and destabilizes the enemy (Figure 5-1). The concept is characterized by its dynamic interoperability, scalability, and



Figure 5-1. The Army's "Robotic and Autonomous Systems Strategy" Foresees Soldier/Robot Teaming and Robot/Robot Teaming as Strategic Mission Enablers (Source: U.S. Army).

integration flexibility, leading to agile integration of new weapons technologies and cost-effective systems [173]. Use of AI is a key enabler for mosaic warfare.

Using AI, the next generation of weapons systems will feature great levels of integrated autonomous functions and decision making. Critical development areas will target AI for greater levels of cross-domain information fusion, intelligent adaptive topology networks, fast data processing, partial information resilient control, decentralized control, and accurate predictive capabilities. Unmanned weapons will feature advanced vision, radar, and sonar algorithms for ATR, multimodal sensor fusion for enhanced perception and intelligent mission planning.

Combat operation is foreseen to engage mission command chain in fast-paced, data-massive operations that human commanders will unlikely cope with. Neurosymbolic AI integrated in the combat management systems is expected to not only unravel the enemy's combat strategies but provide decision support and tactical guidance and disseminate corrective measures in human-understandable form.

Additionally, cross-domain cooperation supported by timely, critical, accurate, and distilled dissemination of combat-valued information constitutes a strategic operational and tactical advantage. Future AI through actionable intelligence is poised to provide the capability of transforming massive situational and battlefield intelligence data (SIGnal INTelligence [SIGINT], Electronic INTelligence [EINT], and Measurement and Signature INTelligence [MASINT]) into coarse intelligent planning that will aid in not only the reduction of cognitive load but in the strategic decision making of commanders and operators. Core functional requirements for actionable intelligence are perception, information fusion, opponent modeling, and intelligent planning.

SECTION 06

AI AND FOREIGN THREATS

Russian and Chinese military and intelligence strategists share the consensus that AI-enabled weapon systems provide a competitive edge and that the adoption of AI in the modern warfare must be a priority for these two nations.

According to a 2018 DoD AI strategy summary report, nations like China and Russia are making important advances and investments in AI for military applications. The investments threaten to erode U.S. technological and operational advantages and destabilize the free and open international order [174]. This section briefly reviews the recent history and state of AI for weapons systems by Russia and China.

6.1 Russia

Russian President Vladimir Putin asserted in September of 2017 that nations leading the research and exploitation of AI will come to dominate global affairs. In December of 2020, Vladimir Putin presented key policies for statewide adoption of AI, one of which is boosting the private investment in domestic AI industries [175].

Although Russia places its AI readiness at 29th in the world, it has advanced from 21st to 16th place in AI research. Russia also ranks 6th in terms of government strategy in an independent global AI index [6].

Russia's global AI initiative via the AI federal project is prioritizing the following areas for more

research and development funding: decision-making systems, computer vision, natural language processing, and advanced robotics and AI technologies. These initiatives, coupled with the military's expression of urgent AI-enabled countermeasures and defensive strategies, are spearheading the Russian military AI development.

Although full autonomy of weapons systems is yet to be fielded, Russia already boasts new semiautonomous glide missiles with "fire and forget" capabilities and semiautonomous nuclear weapons [176]:

- Burevestnik: An experimental, nuclear-powered cruise missile with AI-enabled guidance in development.
- Avangard: An operational, ballistic-launched, hypersonic glide vehicle that computes its glide path before its launch, using an AI-enhanced system, with no explicit predictability to the path it has computed and decided to pursue.
- Poseidon: An experimental, nuclear-powered, autonomous unmanned underwater vehicle (UUV) with complex navigation algorithms and path planning.
- Sarmat: An experimental, intercontinental ballistic missile with AI-enabled guidance.

Swarm technology is another active area of development by the Russian Ministry of Defense, particularly in the naval domain. An antisubmarine

project is underway that seeks to develop AI-enabled robotic swarms with semiautonomous strike capability for swarm detection and submarine searches.

Regarding command and control, Russia's National Defense Management Center and Automated Control System are equipped with AI decision support for real-time and forward-looking analysis of armed conflict and countermeasures networked with other AI-enabled critical nodes and early warning and air-defense systems.

Via AI technologies, the PANTSIR-S air defense systems demonstrated a semiautonomous operation of context-driven target positioning, identification, threat-based prioritization, and optimal solution set for target kills [176].

State-owned Russian media Zvezda released a video of a series of tests on January 23, 2022, at Russia's Alabino training ground, where military drills were conducted to test new robots—logistics unmanned ground vehicles and Scout-Kamikaze robots [177].

While Russia's initiatives on AI and autonomy technology implementation within the military context are ramping up, development efforts are hampered by the absence of digitized infrastructure and a clear strategy similar to U.S. DoD 3000.09 [64].

6.2 China

Unlike Russia, China's ambitions to become an AI superpower can be readily seen through its AI investment and prolific research, the largest in the world. According to Stanford University's Artificial Intelligence Index Report, China has had the largest share of AI research publications in the world since 2017 [175]. The World Intellectual Property Organization (WIPO) reported that China ranked 1st with a 74.7% of the global share of AI-related IP in the past decade and 12th in its 2021 Global Innovation Index, moving up 2 rankings from 2020 [178].

The Chinese military initiatives on AI and autonomous systems are motivated by the need to offset the decades-long conventional superiority of U.S. military technology by shifting and prioritizing its development efforts into disruptive technologies and AI-enabled weapons systems, creating an asymmetric warfare landscape against the United States.

China's AI strategy document provides a view of China's intent and engagement to "promote all kinds of AI technology to become quickly embedded in the field of national defense innovation" [179]. In fact, China is actively engaged in the modernization and "intelligentization" of its military capability across domains of combat operations. U.S. defense analysts indicated that China's Strategic Support Force is pursuing the application of recent advances in AI to its missions of space, cyber, electronic, and psychological warfare [180].

A 2018 assessment of research publications and technical papers from China highlighted a shift in the Chinese military stance from defensive countermeasures to AI-enabled offensive initiatives [181]. The survey also suggested that Chinese experts are pursuing the development of AI-based controls of hypersonic glide vehicles.

Besides the existing platforms, such as the Large Displacement Unmanned Underwater Vehicle HSU-001, Chinese People's Liberation Navy is actively pursuing the development of several unmanned surface vessels (USVs) that may operate semiautonomously [180]. Moreover, the Chinese People's Liberation Army Air Force (PLAAF) is developing advanced UCASs that can optionally operate in MUM-T and unmanned combat settings. An example of these advanced UCASs is the stealth Sharp Sword drone GJ-11 projected to launch swarming decoys at enemy warships [182]. Swarm operations are central to China's strategy in conducting close-in reconnaissance, deception, distributed, and coordinated attacks and saturation attacks [183].

Currently, China's military is operating a Strike Unmanned Combat Air System GJ2, a platform capable of autonomous operations, including reconnaissance, enemy identification, threat judgement, and precision strike [180]. China also operates the weaponized Blowfish A2 helicopter drone with autonomous engagement and strike capability that can accomplish its mission individually or as part of a swarm without a human operator [179].

The PLAAF developed an AI-based autonomic aerial combat simulator used for pilot training and mock within visual-range aerial combat. In a 2021 simulated aerial battle reminiscent of U.S. DARPA's AlphaDogfight competition, the Global Times media outlet reported that an AI pilot shot down seasoned pilots [184]. China intends on integrating the AI pilot into warplanes to assist human pilots during aerial combat.

In its 2021 report, the NSCAI set off an alarm on the U.S. current advancement in AI, alluding to the possibility that the current initiative and state of developments are leaving the United State at a strategic disadvantage [1]. It adds that "America is not prepared to defend or compete in the AI era."

As a result, the DoD has recognized that modern battlefield has transformed into a highly complex environment characterized by growing capabilities of adversaries. Besides the continued development of key technologies in offensive, deterrent, and defensive weapons, the DoD is further seeking to increase weapons efficacy and exploit AI for the symbiosis of humans and intelligent systems on the battlefield. Albeit the transformative effort required to adopt AI across operations, domain enterprises, and services is monumental, the pivotal change is vital for U.S. national security and that of its allies.

This Page Intentionally Left Blank

SECTION 07

ETHICAL CONSIDERATIONS

The use of AI for weapons systems produces a variety of ethical concerns that must be considered. In this section, we will briefly summarize the ethical concerns surrounding AI for weapons systems and the current related U.S. government guidelines.

The overall success of ML and AI has spurred a revolution in data-driven businesses and technologies that touches most industries, not the least of which is the defense sector. Commissioner Nichelle Bechelet has advised member states of the United Nations put a moratorium on the use of AI systems until their adverse risks are addressed [185]. Green [186] laid out 16 major ethical challenges inherent in the automation of interactions between AI and humans. These challenges include psychological, environmental, economic, lack of transparency, interpretability, safety, intent, etc.

The list of ethical implications for use of AI in general is quite extensive, let alone those which apply to the specific use of AI for weapons systems. As such, governments and institutions around the world have begun investigating and formulating strategies that attempt to address these issues and mitigate the risks associated with introducing such technologies.

RAND Corporation's Project Air Force (PAF), a federally funded research and development center, interviewed 29 AI experts to help answer questions regarding the ethical implications of AI in military applications. The experts raised significant concerns about military applications of AI. These

concerns can be grouped into several broad areas—risks of error, increased risks of war, and risks of military operators and leaders putting too much confidence in these capabilities. The experts listed the following concerns of militarized AI that need mitigated: AI systems might make dangerous errors, AI could cause arms racing or escalation, military operators and leaders could put too much trust in AI, and the need for a closer examination of the risks of military AI [187].

An overarching concern among technologists, advocates, and other parties is that military establishments will hasten to integrate AI without paying sufficient regard to the seriousness of these risks. As countries compete to attain the greatest military benefits of AI, they might not put proper precautions in place. Advocates and governments have argued that the key to minimizing most risks is maintaining some level of human agency over these systems. With a human in the AI system's loop, that person can ensure that the system complies with applicable laws and rules of engagement and can be held accountable for the system's actions if it does not.

DoD Directive 3000.09 [64] established DoD policy regarding autonomous and semiautonomous weapons systems and guidelines to minimize its attendant ethical and practical risks. This policy extends to autonomous and semiautonomous, manned and unmanned, and lethal and nonlethal systems. It is important to note that this directive institutes policies and reviews and assigns responsibilities for its own implementation

to appropriate DoD stakeholders. It does not technically allow or disallow any specific practice in AI or autonomy but sets policies where systems be designed to allow commanders and operators exercise “appropriate levels of human judgment over the use of force.” Other policies include, but are not limited to, ensuring that systems seek human input if unable to complete a mission within the commander’s intent or time allotted, rigorous verification and validation and operational test and evaluation, and measures to prevent loss of control of the system to unauthorized parties. However, systems that do not abide by these policies are not automatically disallowed but referred for review “by the Under Secretary of Defense for Policy (USD(P)); the Under Secretary of Defense for Acquisition, Technology, and Logistics (USD(AT&L)); and the CJCS before formal development and again before fielding” [64].

On January 24, 2020, the DoD officially adopted five ethical principles to address concerns surrounding the use of AI in government applications. The listed principles are a commitment from the DoD to provide oversight during the conception and operation of an AI system and ensure responsible use of AI and the protection of privacy and civil liberties. These principles apply to AI systems that serve both combat and noncombat functions [188]:

1. Responsible – DoD personnel remain responsible for the entire AI life cycle.
2. Equitable – DoD will take steps to minimize unintended bias in AI.
3. Traceable – Relevant personnel will have an adequate understanding of the AI system life cycle.
4. Reliable – The DoD’s AI capabilities will have well-defined uses that will undergo testing to ensure safety, security, and effectiveness.
5. Governable – The DoD will have the ability to detect and avoid unintended consequences of AI and can disengage or deactivate deployed AI that demonstrate unintended behavior.

A set of guidelines, known as the Responsible AI (RAI) Guidelines [189], was developed by the Defense Innovation Unit (DIU) to operationalize the DoD’s ethical principles. This was achieved by integrating the ethical principles into the planning, development, and deployment phases of the AI system life cycle. The RAI Guidelines present each phase with a series of questions within a workflow. Progression to each subsequent phase requires all questions are satisfied by program personnel. Projects that followed these guidelines were said to result in functionally superior AI systems that better aligned with the DoD’s ethical principles. Figures 7-1–7-3 show the workflow for the planning, development, and deployment phase of the AI lifecycle, respectively. Questions are tied to relevant, ethical principles adopted by the DoD.

The U.S. government continues to discuss and codify the way in which it will use AI for defense applications. Debates over the ethics of AI for weapons systems will continue, and the U.S. government will be in the spotlight to adopt and implement prudent policies that maintain ethical standards while also enabling all the technological advancements that AI promises.

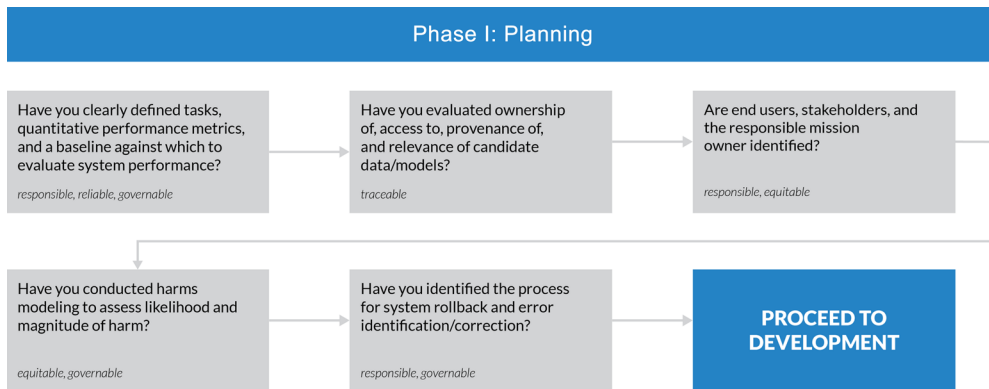


Figure 7-1. The DIU's RAI Guidelines Workflow During the Planning Phase of the AI Life Cycle (Source: Defense Innovation Unit, <https://www.diu.mil/responsible-ai-guidelines> [189]).

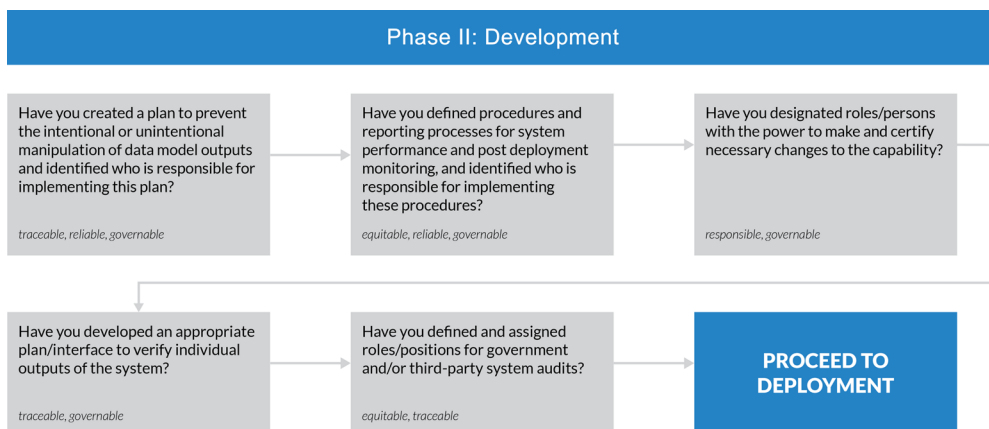


Figure 7-2. The DIU's RAI Guidelines Workflow During the Development Phase of the AI Life Cycle (Source: Defense Innovation Unit, <https://www.diu.mil/responsible-ai-guidelines> [189]).

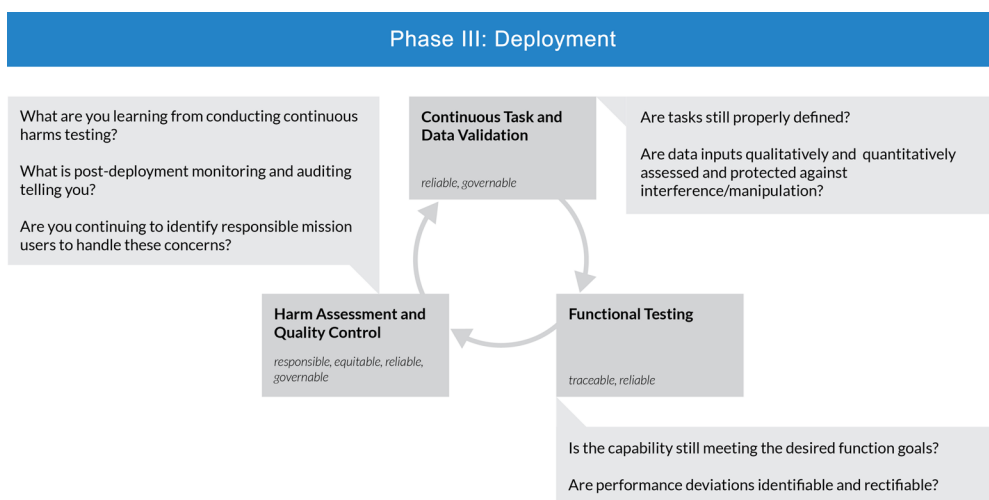


Figure 7-3. The DIU's RAI Guidelines Workflow During the Deployment Phase of the AI Life Cycle (Source: Defense Innovation Unit, <https://www.diu.mil/responsible-ai-guidelines> [189]).

This Page Intentionally Left Blank

SECTION 08

SUMMARY

In view of the advent of the intelligence age and the ever-increasing foreign defiance to the U.S. technological and military leadership, AI is poised to become the determinant factor in ensuring future battlefield readiness and tactical advantage. As combat complexity and information dimensionality increase, it is of prime importance in accordance with the U.S. AI initiatives to identify superlative capabilities that AI technologies can provide to ensure increased combat effectiveness, lethality, and survivability of U.S. weapon systems.

Swarm Intelligence, human-robot interaction, nonlinear adaptive control, visual perception, opponent modeling, and cognitive electronic warfare are some of the key critical areas that will leverage AI to provide bleeding-edge, intelligent combat systems.

Furthermore, with the recent advances in neurosymbolic AI and neuroevolutionary AI research, boundaries are being traversed to close the gap between narrow and general AI. Warfare will never be the same, as AI and related technologies continue to shape combat operations and increase the capabilities and effectiveness of weapons systems across domains and services.

This Page Intentionally Left Blank

REFERENCES

1. National Security Commission on Artificial Intelligence (NCSAI). "National Security Commission on Artificial Intelligence Report." NCSAI, Washington, DC, 2021.
2. Ray, S. "History of AI." *Towards Data Science*, <https://towardsdatascience.com/history-of-ai-484a86fc16ef>, 11 August 2018.
3. Haykin, S. *Neural Networks: A Comprehensive Foundation*. Second edition, Upper Saddle River, NJ: Prentice Hall, 1998.
4. Heaven, W. D. "OpenAI's New Language Generator GPT-3 Is Shockingly Good—and Completely Mindless." *MIT Technology Review*, <https://www.technologyreview.com/2020/07/20/1005454/openai-machine-learning-language-generator-gpt-3-nlp/>, 20 July 2020.
5. Barr, A., and E. Feigenbaum. *The Handbook of Artificial Intelligence: Volume 1*. Butterworth-Heinemann, Stanford, CA: HeurisTech Press, 1981.
6. Russell, S., and P. Norvig. *Artificial Intelligence: A Modern Approach*. Fourth edition, Upper Saddle River, NJ: Pearson, 2020.
7. Wang, P. "On Defining Artificial Intelligence." *Journal of Artificial General Intelligence*, vol. 10, 12 August 2019.
8. Chapelle, O., B. Schölkopf, and A. Zien (editors). *Semi-Supervised Learning*. Cambridge, MA: MIT Press, 2006.
9. Sutton, R., and A. Barto. *Reinforcement Learning: An Introduction*. Second edition, Bradford Books, Cambridge, MA: MIT Press, 2018.
10. Bengio, Y. "Deep Learning of Representations for Unsupervised and Transfer Learning." Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop, vol. 27, pp. 17–37, Bellevue, WA, 2011.
11. Basodi, S., J. Chunyan, H. Zhang, and Y. Pan. "Gradient Amplification: An Efficient Way to Train Deep Neural Networks." *IEEE: Big Data Mining and Analytics*, vol. 3, no. 3, pp. 196–207, 3 September 2020.
12. Goodflew, I., Y. Benio, and A. Courville. *Deep Learning*, p. 282, Cambridge, MA: MIT Press, 2016.
13. Rumelhart, D., G. Hinton, and R. Williams. "Learning Representations by Back-Propagating Errors." *Nature*, vol. 323, pp. 533–536, 1986.
14. Hochreiter, S., and J. Schmidhuber. "LSTM Can Solve Hard Long Time Lag Problems." *Advances in Neural Information Processing Systems 9: Proceedings of the 1997 Conference*, pp. 473–479, Denver, CO, 1997.
15. Kramer, M. A. "Nonlinear Principal Component Analysis Using Autoassociative Neural Networks." *American Institute of Chemical Engineers Journals*, vol. 37, pp. 233–243, 2 February 1991.
16. Billings, S. A. *Nonlinear System Identification: NARMAX Methods in the Time, Frequency, and Spatio-Temporal Domains*. Hoboken, NJ: Wiley, 2013.
17. Kyunghyun, C., B. van Merriënboër, D. Bahdanau, and Y. Bengio. "On the Properties of Neural Machine Translation: Encoder-Decoder Approaches." arXiv:1409.1259, Cornell University, 7 October 2014.
18. He, K., X. Zhang, S. Ren, and J. Sun. "Deep Residual Learning for Image Recognition." 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778, Las Vegas, NV, 2016.
19. Han, S., F. Leng, and Z. Jin. "Speech Emotion Recognition With a ResNet-CNN-Transformer Parallel Neural Network." Proceedings of the 2021 International Conference on Communications, Information System, and Computer Engineering (CISCE), pp. 803–807, Beijing, China, 2021.
20. Huang, G., Z. Liu, L. van der Matten, and K. Weinberger. "Densely Connected Convolutional Networks." 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2261–2269, Honolulu, HI, 2017.
21. Reveriano, F. "The Advantages of DenseNet." *Medium*, <https://medium.com/the-advantages-of-densenet/the-advantages-of-densenet-98de3019cdac>, 24 August 2019.
22. Liu, W., and K. Zeng. "SparseNet: A Sparse DenseNet for Image Classification." arXiv:1804.05340, Cornell University, 15 April 2018.
23. Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. "Generative Adversarial Nets." NIPS' 14: Proceedings of the 17th Neural International Conference on Information Processing Systems, vol. 2, pp. 2672–2680, Montréal, Quebec, Canada, 2014.
24. Sanchez-Lengeling, B., E. Reif, A. Pearce, and A. B. Wiltschko. "A Gentle Introduction to Graph Neural Networks." *Distill*, <https://distill.pub/2021/gnn-intro/>, 2021.
25. Scarselli, F., M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. "The Graph Neural Network Model." *IEEE Transactions on Neural Networks*, vol. 20, pp. 61–80, 1 January 2009.
26. Hinton, G. E., S. Osindero, and Y.-W. Teh. "A Fast Learning Algorithm for Deep Belief Nets." *Neural Computation*, vol. 18, no. 7, pp. 1527–1554, 2006.

REFERENCES, continued

27. Gang, L., L. Xiao, and C. Xiong. "Image Classification With Deep Belief Networks and Improved Gradient Descent." 2017 IEEE International Conference on Computational Science and Engineering (CSE) and IEEE International Conference of Embedded and Ubiquitous Computing (EUC 2017), pp. 375–380, Guangzhou, China, 2017.
28. Schölkopf, B., J. Platt, and T. Hofmann. "Modeling Human Motion Using Binary Latent Variables." *Advances in Neural Information Processing Systems*, pp. 1345–1352, Cambridge, MA: MIT Press, 2007.
29. Voulodimos, A., N. Doulamis, A. Doulamis, and E. Protopapadakis. "Deep Learning for Computer Vision: A Brief Review." *Computer Intelligence and Neuroscience*, vol. 2018, 2018.
30. Kochenderfer, M. J., T. A. Wheeler, and K. H. Wray. *Algorithms for Decision Making*. Cambridge, MA: MIT Press, 2022.
31. Buşoniu, L., R. Babuška, and B. De Schutter. "A Comprehensive Survey of Multiagent Reinforcement Learning." *IEEE Transactions on Systems, Man, and Cybernetics—Part C*, vol. 38, no. 2, pp. 156–172, 2008.
32. Kober, J., J. A. Bagnell, and J. Peters. "Reinforcement Learning in Robotics: A Survey." *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
33. Buşoniu, L., T. de Bruin, D. Tolić, J. Kober, and I. Palunko. "Reinforcement Learning for Control: Performance, Stability, and Deep Approximators." *Annual Reviews in Control*, vol. 46, pp. 8–28, 2018.
34. Wang, S., W. Chaovalitwongse, and R. Babuška. "Machine Learning Algorithms in Bipedal Robot Control." *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 42, no. 5, pp. 728–741, 2012.
35. Zhang, K., Z. Yang, and T. Başar. "Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms." arXiv:1911.10635, Cornell University, 2019.
36. Buşoniu, L., R. Babuška, and B. De Schutter. "Multi-Agent Reinforcement Learning: An Overview." *Innovations in Multi-Agent Systems and Applications—1*, pp. 183–221, Berlin, Germany: Springer, 2010.
37. Nicholson, C. "A Beginner's Guide to Deep Reinforcement Learning." *AI Wiki*, <https://wiki.pathmind.com/deep-reinforcement-learning>, Pathmind, 2022.
38. Lazaridis, A., A. Fachantidis, and I. Vlahavas. "Deep Reinforcement Learning: A State-of-the-Art Walkthrough." *Journal of Artificial Intelligence Research*, vol. 69, pp. 421–1471, September 2020.
39. Petron, A., J. Schneider, B. McGrew, M. Andrychowicz, J. Tobin, S. Sidor, R. Józefowicz, B. Baker, W. Zaremba, M. Plappert, G. Powell, J. Pachocki, M. Chociej, P. Welinder, L. Weng, A. Petron, and A. Ray. "Learning Dexterity." <https://openai.com/blog/learning-dexterity/>, 30 July 2018.
40. Yang, W., B. Guo, Y. Shang, and S. Huang. "Reconfigurable Network Topology Based on Deep Reinforcement Learning in Software-Defined Data-Center Networks." Asia Communications and Photonics Conference/International Conference on Information Photonics and Optical Communications (ACP/IPOC) 2020, pp. 1–3, Beijing, China, 2020.
41. Hannah, L. "Stochastic Optimization." <http://www.stat.columbia.edu/~liam/teaching/compstat-spr14/lauren-notes.pdf>, 4 April 2014.
42. Monteiro, A. C. B., R. P. França, R. Arthur, and Y. Iano. "The Fundamentals and Potential of Heuristics and Metaheuristics for Multiobjective Combinatorial Optimization Problems and Solution Methods." *Multi-Objective Combinatorial Optimization Problems and Solution Methods*, edited by M. Toloo, S. Talatahari, and I. Rahimi, Cambridge, MA: Academic Press, 2022.
43. Blum, C., and A. Roli. "Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison." *ACM Computing Surveys*, vol. 35, no. 3, pp. 268–308, 2003.
44. Miettinen, K. *Nonlinear Multiobjective Optimization*. New York, NY: Springer, 1998.
45. Schranz, M., M. Umlauft, M. Sende, and W. Elmenreich. "Swarm Robotic Behaviors and Current Applications." *Frontiers in Robotics and AI*, vol. 7, <https://doi.org/10.3389/frobt.2020.00036>, 2020.
46. Dhanaraj, N., N. Hewitt, C. Edmonds-Estes, R. Jarman, J. Seo, H. Gunner, A. Hatfield, T. Johnson, L. Yifru, J. Maffeo, G. Pereira, J. Gross, and Y. Gu. "Adaptable Platform for Interactive Swarm Robotics (APIS): A Human Swarm Interaction Research Testbed." The 19th International Conference on Advanced Robotics (ICAR), pp. 720–726, Belo Horizonte, Brazil, 2019.
47. Cunha, A. A. "Swarm Intelligence in Strategy Games." *Computer Science*, <https://fenix.tecnico.ulisboa.pt/downloadFile/281870113702279/ExtendedAbstract.pdf>, May 2015.
48. Telikani, A., A. Tahmassebi, W. Banzhaf, and A. H. Gandomi. "Evolutionary Machine Learning: A Survey." *ACM Computing Surveys*, vol. 54, no. 8, pp. 1–35, November 2022.

REFERENCES, continued

49. Fleming, P., and A. Chipperfield. "An Overview of Evolutionary Algorithms for Control Systems Engineering." <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.101.1923&rep=rep1&type=pdf>, May 2015.
50. Xie, H., Y. Qu, G. Fan, and X. Zhu. "Three-Dimensional Path Planning of UAV Based on Improved Artificial Potential Field." 40th Chinese Control Conference, pp. 7862–7867, Shanghai, China, 2021.
51. Feng, Y., Y. Wu, H. Cao, and J. Sun. "UAV Formation and Obstacle Avoidance Based on Improved APF." 10th International Conference on Modelling, Identification, and Control, Guiyang, China, 2018.
52. Fan, X., Y. Guo, H. Liu, B. Wei, and W. Lyu. "Improved Artificial Potential Field Method Applied for AUV Path Planning." *Mathematical Problems in Engineering*, vol. 2020, 2020.
53. Batista, J., D. Souza, J. L. Nunes de Silva, J. Rodrigues da Costa, L. L. N. R. Laurinda, K. Ramos, and A. Braga. "Trajectory Planning Using Artificial Potential Fields With Metaheuristics." *IEEE Latin America Transactions*, vol. 18, no. 5, pp. 914–922, May 2020.
54. Biswas, A., K.K. Mishra, S. Tiwari, and A. K. Misra. "Physics-Inspired Optimization Algorithms: A Survey." *Journal of Optimization*, vol. 2013, 12 June 2013.
55. Shi, J., Q. Zhang, and E. Tsang. "EB-GLS: An Improved Guided Local Search Based on the Big Valley Structure." arXiv:1709.07576[cs.AI], Cornell University, 22 September 2017.
56. Glover, F., M. Laguna, and R. Martí. "Principles and Strategies of Tabu Search." *Handbook of Approximation Algorithms and Metaheuristics*, second edition, edited by T. F. Gonzalez, New York, NY: Chapman and Hall/CRC Press, 2018.
57. García-Martínez, C., and M. Lozano. "Local Search Based on Genetic Algorithms." *Advances in Metaheuristics for Hard Optimization*, https://sci2s.ugr.es/sites/default/files/ficherosPublicaciones/0609_garciaSpringerBook08.pdf, 2007.
58. Karthik, K., N. Sharma, C. Dharmatti, and J. E. Siegel. "A Survey of Path Planning Algorithms for Mobile Robots." *MDPI: Vehicles*, vol. 3, pp. 448–468, 4 August 2021.
59. Dickson, B. "Neuro-Symbolic AI Brings Us Closer to Machines With Common Sense." TechTalks Newsletter, <https://bdtechtalks.com/2022/03/14/neuro-symbolic-ai-common-sense/>, 14 March 2022.
60. Susskind, Z., B. Arden, and L. K. John. "Neuro-Symbolic AI: An Emerging Class of AI Workloads and Their Characterization." arXiv:2109.06133, Cornell University, 13 September 2021.
61. Stanley, K. O., and M. Risto. "Evolving Neural Networks Through Augmenting Topologies." *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, 2002.
62. Stanley, K. O., and J. Cune. "Designing Neural Networks Through Neuroevolution." *Nature Machine Intelligence*, vol. 1, pp. 24–35, 7 January 2019.
63. Pratihari, D. K., and L. C. Jain. *Intelligent Autonomous Systems: Foundations and Applications*. Springer Verlag, Berlin, 2010.
64. Deputy Secretary of Defense. "Autonomy in Weapon Systems." DoD Directive 3000.09, <http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf>, U.S. Department of Defense, Washington, DC, 21 November 2012.
65. Ilchinski, A. "AI, Robots, and Swarms: Issues, Questions, and Recommended Studies." CNA Analysis and Solutions, Arlington, VA, January 2017.
66. U.S. Department of Defense, Defense Science Board. "The Role of Autonomy in DoD Systems." Office of the Under Secretary of Defense for Acquisition, Technology, Logistics, and Defense, Washington, DC, July 2012.
67. Beer, J. M., A. D. Fisk, and W. A. Rogers. "Toward a Framework for Levels of Robot Autonomy in Human-Robot Interaction." *Journal of Human-Robot Interaction*, vol. 3, no. 2, pp. 74–99, 10 July 2014.
68. Huany, H.-M. (editor). "Autonomy Levels for Unmanned Systems (ALFUS) Framework—Volume I: Terminology (Version: 2)." NIST SP 1011-I-2.0, National Institute of Standards and Technology, Gaithersburg, MD, 2008.
69. U.S. Department of Defense, Defense Science Board. "Autonomy Report." U.S. Department of Defense, Washington, DC, 2012.
70. Finn, A., and S. Scheding. "Developments and Challenges for Autonomous Unmanned Vehicles." *Intelligent Systems Reference Library*, vol. 3, Springer-Verlag, Berlin, Heidelberg, 2010.
71. Barla, N. "The Beginner's Guide to Semantic Segmentation." <https://www.v7labs.com/blog/semantic-segmentation-guide>, V7Labs, May 2022.
72. Bandyopadhyay, H. "The Definitive Guide to Instance Segmentation." <https://www.v7labs.com/blog/instance-segmentation-guide>, V7Labs, July 2022.
73. Blasch, E. P., and Y. Zheng. "Multi-Modal Video Fusion for Context-Aided Tracking." IEEE 23rd International Conference on Information Fusion (FUSION), pp. 1–8, Pilanesberg National Park, South Africa, July 2020.

REFERENCES, continued

74. Huang, Z., C. Lv, Y. Xing, and J. Wu. "Multi-Modal Sensor Fusion-Based Deep Neural Network for End-to-End Autonomous Driving With Scene Understanding." *IEEE Sensors Journal*, vol. 21, no. 10, pp. 11781–11790, 15 May 2021.
75. Abbeel, P. "Apprenticeship Learning and Reinforcement With Application to Robotic Control." Doctoral dissertation, Stanford University, Stanford, CA, August 2008.
76. Khalid, M., and S. Omatu. "A Neural Network Based Control Scheme With an Adaptive Neural Model Reference Structure." *IEEE International Joint Conference on Neural Networks*, vol. 3, pp. 2128–2133, Singapore, November 1991.
77. Khalil, H. K. *Nonlinear Systems*. Third edition, Upper Saddle River, NJ: Pearson, 18 December 2001.
78. Bensoussan, A., Y. Li, D. P. C. Nguyen, M.-B. Tran, S. C. P. Yim, and X. Zhou. "Machine Learning and Control Theory." arXiv:2006.05604, Cornell University, 10 June 2020.
79. Shi, G., K. Azizzadenesheli, M. O'Connell, S.-J. Chung, and Y. Yue. "Meta-Adaptive Nonlinear Control: Theory and Algorithms." arXiv:2106.06098v3, Cornell University, 26 October 2021.
80. Shi, G., X. Shi, M. O'Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung. "Neural Lander: Stable Drone Landing Control Using Learned Dynamics." arXiv:1811.08027, Cornell University, 2019.
81. Lin, C., F. Jiang, Z. Wang, and J. Li. "Multiconstrained Real-Time Entry Guidance Using Deep Neural Networks." *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 1, pp. 25–340, February 2021.
82. Du, Y., D. Sun, J. Fu, and P. Zhang. "Adaptive Maneuver Control of Hypersonic Re-Entry Flight via Self Organizing Recurrent Functional Link Network." International Joint Conference on Neural Networks (IJCNN), Vancouver, British Columbia, Canada, 2016.
83. Takahashi, K., L. Wantanabe, H. Yamasaki, S. Hiraoka, and M. Hashimoto. "Remarks on Control of a Robot Manipulator Using a Quaternion Recurrent Neural-Network-Based Compensator." Australian and New Zealand Control Conference (ANZCC), pp. 244–247, Gold Coast, Australia, November 2020.
84. Kumar, A., Z. Fu, D. Pathak, and J. Malik. "RMA: Rapid Motor Adaptation for Legged Robots." arXiv:2107.04034, Cornell University, 2021.
85. Dooraki, A. R., and D.-J. Lee. "An Innovative Bio-Inspired Flight Controller for Quad-Rotor Drones: Quad-Rotor Drone Learning to Fly Using Reinforcement Learning." *Robotics and Autonomous Systems*, vol. 135, January 2021.
86. Chithapuram, C., Y. V. Jeppu, and C. A. Kumar. "Artificial Intelligence Guidance for Unmanned Aerial Vehicles in Three-Dimensional Space." International Conference on Contemporary Computing and Informatics (IC3I), pp. 1256-1261, Mysore, India, November 2014.
87. Wang, G., A. Hao, C. Yu, and G. Ziyi. "Deep Reinforcement Learning-Based Backstepping Control of Air-Breathing Hypersonic Vehicles With Actuator Constraints." The 40th Chinese Control Conference, Shanghai, China, 2021.
88. Furfaro, R., A. Scorsoglio, R. Linares, and M. Massari. "Adaptive Generalized ZEM-ZEV Feedback Guidance for Planetary Landing via a Deep Reinforcement Learning Approach." *Acta Astronautica*, vol. 171, pp. 156–171, 2020.
89. Iwasaki, H., and A. Okuyama. "Development of a Reference Signal Self-Organizing Control System Based on Deep Reinforcement Learning." IEEE International Conference on Mechatronics (ICM), pp. 1–5, Kashiwa, Japan, March 2021.
90. Abbeel, P. "Apprenticeship Learning and Reinforcement." Doctoral dissertation, Stanford University, Stanford, CA, August 2008.
91. Couto, G. C. K., and E. A. Antonelo. "Generative Adversarial Imitation Learning for End-to-End Autonomous Driving on Urban Environments." IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1–7, Orlando, FL, 2021.
92. Merel, J., Y. Tassa, D. TB, S. Srinivasan, J. Lemmon, Z. Wang, G. Wayne, and N. Heess. "Learning Human Behaviors From Motion Capture by Adversarial Imitation." arXiv:1707.02201v2, Cornell University, 2017.
93. Lucia, S., and F. Fiedler. "Model Predictive Control Python Toolbox." Robust Optimal Control Toolbox, revision 115676db, https://www.do-mpc.com/en/v4.1.1/theory_mpc.html, 22 March 2022.
94. Lenz, I., R. Knepper, and A. Saxena. "DeepMPC: Learning Deep Latent Features for Model Predictive Control." <http://deepmpc.cs.cornell.edu/DeepMPC.pdf>, Department of Computer Science, Cornell University, 24 March 2022.
95. Bieker, K., S. Peitz, S. L. Brunton, J. N. Kutz, and M. Dellnitz. "Deep Model Predictive Control With Online Learning for Complex Physical Systems." arXiv:1905.10094, Cornell University, 24 May 2019.

REFERENCES, continued

96. Mur-Artal, R., J. M. M. Montiel, and J. D. Tardos. "ORB-SLAM: A Versatile and Accurate Monocular SLAM System." arXiv:1502.00956, Cornell University, 3 February 2015.
97. Alejo, C., and J. Civera. "DPPTAM: Dense Piecewise Planar Tracking and Mapping from a Monocular Sequence." IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5686–5693, Hamburg, Germany, September 2015.
98. Jiang, M., S. Song, Y. Li, and W. Jin. "A Survey of Underwater Acoustic SLAM System." Intelligent Robotics and Applications, Conference Proceedings, pp. 159–170, Shenyang, China, August 2019.
99. Shen, C., X. Guo, X. Chen, H. Cao, J. Li, J. Liu, J. Tang, and Y. Zhang. "Seamless GPS/Inertial Navigation System Based on Self-Learning Square-Root Cubature Kalman Filter." *IEEE Transactions on Industrial Electronics*, vol. 68, no. 1, pp. 499–508, January 2021.
100. Kube, R., D. Brunner, F. M. Bianchi, and B. LaBombard. "Outlier Classification Using Autoencoders: Application for Fluctuation Driven Flows in Fusion Plasmas." *The Review of Scientific Instruments*, 16 January 2019.
101. Sakurada, M., and T. Yairi. "Anomaly Detection Using Autoencoders With Nonlinear Dimensionality Reduction." Proceedings of the MLSDA 2nd Workshop on Machine Learning for Sensory Data Analysis, pp. 4–11, Gold Coast, Australia, December 2014.
102. Bakarji, J., K. Champion, J. N. Kutz, and S. L. Brunton. "Discovering Governing Equations From Partial Measurements With Deep Delay Autoencoders." arXiv:2201.05136, Cornell University, 2022.
103. Guo, T., T. Lin, and Y. Lu. "An Interpretable LSTM Neural Network for Autoregressive Exogenous Model." arXiv:1804.05251v1, Cornell University, April 2018.
104. Liu, H., and X. Song. "Nonlinear System Identification Based on NARX Network." The 10th Asian Control Conference (ASCC), Kota Kinabalu, Malaysia, June 2015.
105. Hefron, R. "Air Combat Evolution (ACE)." <https://www.darpa.mil/program/air-combat-evolution>, Defense Advanced Research Projects Agency, Arlington, VA, accessed on 12 February 2022.
106. Chen, B., and G. Quan. "NP-Hard Problems of Learning From Examples." Fifth International Conference on Fuzzy Systems and Knowledge Discovery, pp. 182–186, Jinan, China, October 2008.
107. Yu, Z., F. Sun, X. Lu, and Y. Song. "Overview of Research on 3D Path Planning Methods for Rotor UAV." International Conference on Electronics, Circuits, and Information Engineering (ECIE), Zhengzhou, China, April 2021.
108. Montazeri, A., A. Can, and I. H. Imran. "Unmanned Aerial Systems: Autonomy, Cognition, and Control." Unmanned Aerial Systems: Theoretical Foundation and Applications, chapt. 3, pp. 47–80, edited by A. Koubaa and A. T. Azar, Cambridge, MA: Academic Press, 2021.
109. Janét, R., C. Luo, and M. G. Kay. "The Essential Visibility Graph: An Approach to Global Motion Planning for Autonomous Mobile Robots." Proceedings of 1995 IEEE International Conference on Robotics and Automation, vol. 2, pp. 1958–1963, Nagoya, Japan, May 1995.
110. Kusyik, J., M. U. Uyar, K. Ma, J. J. Wu, W. Ruan, D. K. Guha, G. Bertoli, and J. Boksiner. "AI Based Flight Control for Autonomous UAV Swarms." International Conference on Computational Science and Computational Intelligence, Las Vegas, NV, December 2018.
111. Yang, F., P. Wang, Y. Zhang, L. Zheng, and J. Lu. "Survey of Swarm Intelligence Optimization Algorithm." IEEE International Conference on Unmanned Systems, Beijing, China, October 2017.
112. Dan, Z., H. Hongyan, and H. Yu. "The Optimal Selection of the Parameters for the Ant Colony Algorithm With Small-Perturbation." International Conference on Computing, Control, and Industrial Engineering, Wuhan, China, June 2010.
113. Li, P., and H. Zhu. "Parameter Selection for Ant Colony Algorithm Based on Bacterial Foraging Algorithm." *Mathematical Problems in Engineering*, vol. 2016, 2016.
114. Friedrich, T., T. Kötzing, M. S. Krejca, and A. M. Sutton. "Robustness of Ant Colony Optimization to Noise." *Evolutionary Computation*, vol. 24, pp. 237–254, 2016.
115. Konatowski, S., and P. Pawłowski. "Ant Colony Optimization Algorithm for UAV Path Planning." The 2018 14th International Conference on Advanced Trends in Radioelectronics, Telecommunications, and Computer Engineering, Lviv-Slavske, Ukraine, February 2018.
116. Cekmez, U., M. Ozsiginan, and O. K. Sahingoz. "Multi Colony Ant Optimization for UAV Path Planning With Obstacle Avoidance." International Conference on Unmanned Aircraft Systems (ICUAS), Arlington, VA, June 2016.
117. Chen, L., W.-L. Liu, and J. Zhong. "An Efficient Multi-Objective Ant Colony Optimization for Task Allocation of Heterogeneous Unmanned Aerial Vehicles." *Journal of Computational Science*, vol. 58, February 2022.

REFERENCES, continued

118. Wu, H.-S., and F.-M. Zhang. "Wolf Pack Algorithm for Unconstrained Global Optimization." *Mathematical Problems in Engineering*, 2014.
119. Jin Jangang, J., W. Husheng, Z. Renjun, M. Rafik, Z. Xuanwu, W. Hu, J. H. Wu, R. Zhan, R. Menassel, and X. Zhou. "Self-Organized Search-Attack Mission Planning for UAV Swarm Based on Wolf Pack Hunting Behavior." *Journal of Systems Engineering and Electronics*, vol. 32, no. 6, pp. 1463–1476, December 2021.
120. Nguyen, H. T., M. Garratt, L.-T. Bui, and H. A. Abbass. "Apprenticeship Learning for Continuous State Spaces and Actions in a Swarm-Guidance Shepherding Task." IEEE Symposium Series on Computational Intelligence, SSCI 2019, pp. 102–109, Xiamen, China, December 2019.
121. Saied, M., M. Sim, F. Clovis, H. Shraim, and H. Mazeh. "Unmanned Aerial Vehicles Fleet Control via Artificial Bee Colony Algorithm." Fourth International Conference on Control and Fault Tolerant Systems (SysTol), Casablanca, Morocco, 2019.
122. Li, X., M. Jian, X. Zhang, and S. Jia. "Design of the Fruit Fly Optimization Algorithm Based Path Planner for UAV in 3D Environments." IEEE International Conference on Mechatronics and Automation (ICMA), pp. 381–386, Takamatsu, Kagawa, Japan, August 2017.
123. Iscan, H., M. S. Kiran, and M. Gunduz. "A Novel Candidate Solution Generation Strategy for Fruit Fly Optimizer." *IEEE Access*, vol. 7, pp. 130901–130921, 2019.
124. Zhu, H., Y. Wang, Z. Ma, and X. Li. "A Comparative Study of Swarm Intelligence Algorithms for UCAV Path-Planning Problems." *Mathematics*, vol. 9, no. 2, 2021.
125. Meng, X. B., Y. Liu, and H. Zheng. "A New Bio-Inspired Algorithm: Chicken Swarm Optimization." International Conference in Swarm Intelligence, Hefei, China, 2014.
126. Liang, X., D. Kou, and L. Wen. "An Improved Chicken Swarm Optimization Algorithm and Its Application in Robot Path Planning." *IEEE Access*, vol. 8, pp. 49543–49550, January 2020.
127. Robson, J. M. "The Complexity of Go." Proceedings of IFIP 9th World Computer Congress, pp. 413–417, Paris, France, September 1983.
128. Fraenkel, A. S., and D. Lichtenstein. "Computing a Perfect Strategy for $n \times n$ Chess Requires Time Exponential in n ." *Journal of Combinatorial Theory*, series A, vol. 31, no. 2, pp. 199–214, 1981.
129. Silver, D., A. Huang, C. Maddison, A. Guez, L. Sifre, G. van der Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Graepel, T. Lillicrap, M. Leach, K. Kavukcuoglu, and D. Hassabis. "Mastering the Game of Go With Deep Neural Networks and Tree Search." *DeepMind*, <https://deepmind.com/research/publications/2019/mastering-game-go-deep-neural-networks-tree-search>, January 2016.
130. Silver, D., J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van der Driessche, T. Graepel, and D. Hassabis. "Mastering the Game of Go Without Human Knowledge." *Nature*, https://www.nature.com/articles/nature24270.epdf?author_access_token=VJXbVjaSHxFOctQQ4p2k4tRgN0jAjWel9jnR3ZoTv0PVW4gB86EEpGqTRDtpIz-2rmo8-KG06gqVobU5NSCFeHILHcVFUeMsbvwS-lxjQGG98faovwjxeTUgZAUMnRQ, 18 October 2017.
131. Schrittwieser, J., I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, T. Lillicrap, and D. Silver. "Mastering Atari, Go, Chess, and Shogi by Planning With a Learned Model." *Nature*, https://www.nature.com/articles/s41586-020-03051-4.epdf?sharing_token=kTk-xTZpQOF8Ym8nTQK6EdRgN0jAjWel9jnR3ZoTv0PMSWGj38iNlyNOw_ooNp2BvzZ4nlcedo7GEXD7UmLqb0M_V_fop31mMY9VBBLNmGbm0K9jETKkZnJ9SgJ8Rwhp3ySvLuTcUr888puiYbngQ0fiMf45ZGDAQ7fUI66-u7Y%3D, 23 December 2020.
132. Yekun, C. "AlphaStar: Grandmaster Level in StarCraft II Explained." <https://ychai.uk/notes/2019/07/21/RL/DRL/Decipher-AlphaStar-on-StarCraft-II/>, 2019.
133. Shek, J. "Takeaways From OpenAI Five." *Towards Data Science*, <https://towardsdatascience.com/takeaways-from-openai-five-2019-f90a612fe5d>, 23 April 2019.
134. Nashed, S. B., and S. Zilberstein. "A Survey of Opponent Modeling in Adversarial Domains." *Journal of Artificial Intelligence Research*, vol. 73, pp. 277–327, January 2022.
135. Tang, Z., Y. Zhu, D. Zhao, and S. M. Lucas. "Enhanced Rolling Horizon Evolution Algorithm With Opponent Model Learning." *IEEE Transactions on Games*, 2022.
136. Synnaeve, G., Z. Lin, J. Gehring, D. Gant, V. Mella, V. Khalidov, N. Carion, and N. Usunier. "Forward Modeling for Partial Observation Strategy Games—A StarCraft Defogger." arXiv:1812.00054, Cornell University, November 2018.
137. Chen, H., J. Huang, Q. Liu, C. Wang, and H. Deng. "Detecting and Tracing Multi-Strategic Agents With Opponent Modelling and Bayesian Policy Reuse. The 15th IEEE Conference on Industrial Electronics and Applications (ICIEA), pp. 1098–1103, Kristiansand, Norway, November 2020.

REFERENCES, continued

138. Chen, H., C. Wang, J. Huang, and J. Gong. "Efficient Use of Heuristics for Accelerating XCS-Based Policy Learning in Markov Games." arXiv:2005.12553v1, Cornell University, 2020.
139. Chen, S., and R. C. Arkin. "Counter-Misdirection in Behavior-Based Multi-Robot Teams." IEEE International Conference on Intelligence and Safety for Robotics (ISR), pp. 268–275, Nagoya, Japan, March 2021.
140. Zheng, Y., Z. Meng, J. Hao, Z. Zhang, T. Yang, and C. Fan. "A Deep Bayesian Policy Reuse Approach Against Non-Stationary Agents." The 32nd International Conference on Neural Information Processing Systems, pp. 962–972, Montréal, Canada, 2018.
141. He, H., J. Boyd-Graber, K. Kwok, and H. Daumé III. "Opponent Modeling in Deep Reinforcement Learning." arXiv:1609.05559, Cornell University, September 2016.
142. Biro, P., and S. G. Walker. "A Reinforcement Learning Based Approach to Play Calling in Football." arXiv:2103.06939v1, Cornell University, March 2021.
143. Ammar, M. A., M. S. Abdel-Latif, H. A. Hassan, and K. M. Badran. "Deep Learning Achievements and Opportunities in Domain of Electronic Warfare Applications." International Conference on Intelligent Computing and Information Systems (ICICIS), Cairo, Egypt, December 2021.
144. Sharma, P., K. K. Sarma, and N. E. Mastorakis. "Artificial Intelligence Aided Electronic Warfare Systems—Recent Trends and Evolving Applications." *IEEE Access*, vol. 8, pp. 224761–224780, 2020.
145. Liao, X., B. Li, and B. Yang. "A Novel Classification and Identification Scheme of Emitter Signals Based on Ward's Clustering and Probabilistic Neural Networks With Correlation Analysis." *Computational Intelligence and Neuroscience*, vol. 2018, June 2018.
146. Qu, Z., X. Mao, and Z. Deng. "Radar Signal Intra-Pulse Modulation Recognition Based on Convolutional Neural Network." *IEEE Access*, vol. 6, pp. 43874–43884, August 2018.
147. O'Shea, T., T. Roy, and T. Clancy. "Learning Robust General Radio Signal Detection Using Computer Vision Methods." The 51st Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, 2017.
148. Lee, G.-H., J. Jo, and C. H. Park. "Jamming Prediction for Radar Signals Using Machine Learning Methods." *Security and Communication Networks*, vol. 2020, 2020.
149. BAE Systems. "DARPA Funding Brings Machine Learning to BAE Systems' Signals Intelligence Capabilities." <https://www.baesystems.com/en-us/article/darpa-funding-brings-machine-learning-to-bae-systems-signals-intelligence-capabilities>, 8 July 2019.
150. Qiang, X., Z. Wei-Gang, and B. Yuan. "Jamming Style Selection for Small Sample Radar Jamming Rule Base." IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC), Qingdao, China, September 2018.
151. Amuru, S., C. Tekin, M. van der Schaar, and R. M. Buehrer. "A Systematic Learning Method for Optimal Jamming." IEEE International Conference on Communications (ICC), London, UK, 2015.
152. Yangyang, L., X. Wang, D. Liu, Q. Guo, J. Zhang, and Y. Xu. "On the Performance of Deep Reinforcement Learning-Based Anti-Jamming Method Confronting Intelligent Jammer." *Applied Sciences*, vol. 9, no. 7, p. 1361, March 2019.
153. Kang, L., and J. Bo. "Reinforcement Learning Based Anti-Jamming Frequency Hopping Strategies Design for Cognitive Radar." IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC), Qingdao, China, September 2018.
154. Albert, L. "AlphaDogfight Trial Winner Heron Systems Acquired by Shield AI." Overt Defense (OVD), <https://www.overtdefense.com/2021/07/23/alphadogfight-trial-winner-heron-systems-acquired-by-shield-ai/>, 23 July 2021.
155. Tingley, B. "New Swarming Capability Planned for the V-Bat Vertical Takeoff and Landing Drone." *The Warzone*, <https://www.thedrive.com/the-war-zone/41755/new-swarming-capability-planned-for-the-v-bat-vertical-takeoff-and-landing-drone>, 29 July 2021.
156. Trevithick, J. "The Fight for the Air Force's 'Skyborg' Artificial Intelligence Equipped Drones Has Begun." *The Warzone*, <https://www.thedrive.com/the-war-zone/33567/the-fight-for-the-air-forces-skyborg-artificial-intelligence-equipped-drones-has-begun>, 22 May 2020.
157. U.S. Air Force Research Laboratory Public Affairs. "AFRL Successfully Completes XQ-58A Valkyrie Flight and Payload Release Test." *Hill Airforce Base News*, <https://www.hill.af.mil/News/Article-Display/Article/2563194/afrl-successfully-completes-xq-58a-valkyrie-flight-and-payload-release-test/>, 5 April 2021.
158. Insinna, V. "Skyborg Makes Its Second Flight, This Time Autonomously Piloting General Atomics' Avenger Drone." C4ISRNET, <https://www.c4isrnet.com/unmanned/2021/06/30/skyborg-makes-its-second-flight-this-time-autonomously-piloting-general-atomics-avenger-drone/>, 30 June 2021.

REFERENCES, continued

159. QinetiQ News. "QinetiQ North America and Pratt and Miller Defense Announce Robotic Combat Vehicle Partnership." <https://www.qinetiq.com/en/news/qinetiq-north-america-and-pratt-and-miller-defense-announce-robotic-combat-vehicle-partnership>, accessed on 22 March 2022.
160. U.S. Air Force Research Laboratory (AFRL). "Golden Horse Colosseum." <https://afresearchlab.com/technology/vanguards/successstories/golden-horde>, accessed on 27 February 2022.
161. Toon, J. "Collaborative Weapons Work Together to Respond to Changing Conditions." Georgia Tech Research Institute (GTRI) 2021 Annual Report, <https://gtri.gatech.edu/newsroom/collaborative-weapons-work-together-respond-changing-conditions>, 11 November 2021.
162. Gibbons-Neff, T. "Watch the Pentagon's New Hive-Mind-Controlled Drone Swarm in Action." *The Washington Post*, <https://www.washingtonpost.com/news/checkpoint/wp/2017/01/10/watch-the-pentagons-new-hive-mind-controlled-drone-swarm-in-action/>, 10 January 2017.
163. Airforce Technology. "Mako Unmanned Tactical Aerial Platform-22 (UTAP22)." <https://www.airforce-technology.com/projects/mako-unmanned-tactical-aerial-platform-22-utap22/>, 17 May 2021.
164. Raytheon, Missiles and Defense. "Coyote UAS." <https://www.raytheonmissilesanddefense.com/what-we-do/counter-uas/effectors/coyote>, 23 March 2022.
165. Wolf, M. T., A. Rahmani, J.-P. de la Croix, G. Woodward, J. Vander Hook, D. Brown, S. Schaffer, C. Lim, P. Bailey, S. Tepsuporn, M. Pomerantz, V. Nguyen, C. Sorice, and M. Sandoval. "CARACaS Multi-Agent Maritime Autonomy for Unmanned Surface Vehicles in the Swarm II Harbor Patrol Demonstration." <https://trs.jpl.nasa.gov/bitstream/handle/2014/47748/CL%2317-1581.pdf>, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 2017.
166. Hudson, A. "The Looming Swarm." *Airforce Magazine*, <https://www.airforcemag.com/article/the-looming-swarm/>, 22 March 2019.
167. DRAPER. "Swarms of Underwater Drones Deliver the Goods in Naval Exercise." <https://www.draper.com/news-releases/swarms-underwater-drones-deliver-goods-naval-exercise>, 13 September 2018.
168. HRL Laboratories. "NEMESIS Keeps an Eye on Warfighters." <https://www.hrl.com/news/2020/03/10/nemesis-keeps-an-eye-on-warfighters>, 2020.
169. IG (Inside GNSS [Global Navigation Satellite Systems] Engineering, Policy, and Design). "GPS Threat Detector Software Scans RF Environment for Jamming & Spoofing Indicators." <https://insidegnss.com/gps-threat-detector-software-scans-rf-environment-for-jamming-spoofing-indicators/>, 28 October 2021.
170. Erwin, S. "Pentagon Report: China Amassing Arsenal of Anti-Satellite Weapons." *Space News*, <https://spacenews.com/pentagon-report-china-amassing-arsenal-of-anti-satellite-weapons/>, 1 September 2020.
171. Keller, J. "Air Force Picks KBR to Use Artificial Intelligence (AI) in Cognitive Position, Navigation, and Timing (PNT)." *Military Aerospace Electronics*, <https://www.militaryaerospace.com/computers/article/14210872/cognitive-artificial-intelligence-ai-position-navigation-and-timing-pnt>, 23 September 2021.
172. Defense Advanced Research Projects Agency. "DARPA Tiles Together a Vision of Mosaic Warfare." <https://www.darpa.mil/work-with-us/darpa-tiles-together-a-vision-of-mosaic-warfare>, 2018.
173. Defense Advanced Research Projects Agency. "Strategic Technology Office Outlines Vision for Mosaic Warfare." <https://www.darpa.mil/news-events/2017-08-04>, 4 August 2017.
174. U.S. Department of Defense. "Summary of the 2018 Department of Defense Artificial Intelligence Strategy, Harnessing AI to Advance Our Security and Prosperity," 2018.
175. Zhang, D., S. Mishra, E. Brynjolfsson, J. Etchemendy, D. Ganguli, B. Grosz, T. Lyons, J. Manyika, J. C. Niebles, M. Sellitto, Y. Shoham, J. Clark, and R. Perrault. "Artificial Intelligence Index Report." AI Index Committee, Human Centered Artificial Intelligence, Stanford University, Stanford, CA, March 2021.
176. Edmonds, J., S. Bendett, A. Fink, M. Chestnut, D. Gorenburg, M. Kofman, K. Stricklin, and J. Waller. "Artificial Intelligence and Autonomy in Russia." Center for Naval Analyses (CNA), Arlington, VA, 1 April 2021.
177. Atherton, K. D. "These Russian Military Robots Help With Logistics and Scouting." <https://www.popsci.com/technology/russian-soldiers-test-new-military-robots/>, 6 February 2022.
178. World Intellectual Property Organization (WIPO). *WIPO News*. https://www.wipo.int/about-wipo/en/offices/china/news/2021/news_0037.html, 28 September 2021.
179. Gibson, A., A. J. Merchant, and B. D. Vigneron. "Autonomous Systems in the Combat Environment: The Key or the Curse to the U.S." *The Strategy Bridge*,

REFERENCES, continued

- <https://thestrategybridge.org/the-bridge/2020/10/8/autonomous-systems-in-the-combat-environment-the-key-or-the-curse-to-the-us>, 8 October 2020.
180. Kania, E. B. "AI Weapons in China's Military Innovation." *Global China, Assessing China's Growing Role in the World*, The Brookings Institution, April 2020.
181. Shazeda, A., N. E. Bajema, S. Bendett, B. A. Chang, R. Creemers, C. C. Demmchak, S. W. Denton, J. Ding, S. Hoffman, R. Joseph, E. Kania, J. Kerr, L. Kostopoulos, J. A. Lewis, M. Libicki, H. Lin, K. Miura, R. Morgus, R. E. Odell, E. Pauwels, L. Saalman, J. Snow, L. Steckman, and V. Weber. "AI, China, Russia, and the Global Order: Technological, Political, Global, and Creative Perspectives." *A Strategic Multilayer Assessment (SMA) Periodic Publication*, p. 153, December 2018.
182. Paleja, A. "China Unveils Long-Range GJ-11 Drones That Can Launch Decoys at Enemy Warships." *Interesting Engineering*, <https://interestingengineering.com/china-unveils-long-range-gj-11-drones-that-can-launch-decoys-at-enemy-warships>, 11 October 2021.
183. Pollpeter, K., and A. Kerrigan. "The PLA and Intelligent Warfare: A Preliminary Analysis." Center for Naval Analyses (CNA), Arlington, VA, 2021.
184. Trevithick, J. "Chinese Pilots Are Also Dueling With AI Opponents in Simulated Dogfights and Losing: Report." *The Warzone*, <https://www.thedrive.com/the-war-zone/41152/chinese-pilots-are-also-dueling-with-ai-opponents-in-simulated-dogfights-and-losing-report>, 18 June 2021.
185. Neuman, S. "The U.N. Warns That AI Can Pose a Threat to Human Rights." *NPR*, 16 September 2021.
186. Green, B. P. "Artificial Intelligence and Ethics: Sixteen Challenges and Opportunities." *All About Ethics*, <https://www.scu.edu/ethics/all-about-ethics/artificial-intelligence-and-ethics-sixteen-challenges-and-opportunities/>, Markkula Center for Applied Ethics, Santa Clara University, 2020.
187. Forrest, M., B. Bourdreaux, A. J. Lohn, M. Ashby, C. Curriden, K. Kilma, and D. Grossman. "Military Applications of Artificial Intelligence: Ethical Concerns in an Uncertain World." <https://apps.dtic.mil/sti/citations/AD1097313>, Defense Technical Information Center, 1 January 2020.
188. U.S. Department of Defense. "DoD Adopts Ethical Principles for Artificial Intelligence." <https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>, DoD Release, 24 February 2020.
189. Dunnmon, J., B. Goodman, P. Kirechu, C. Smith, and A. Deusen. "Responsible AI Guidelines in Practice." https://assets.ctfassets.net/3nanhbfkr0pc/acoo1Fj5uungnGNP/J3QWy/c800b2b32b59cd21b7f2357dd0ce899c/2021_RAI_Report.pdf, Defense Innovation Unit, 2021.

This Page Intentionally Left Blank

BIBLIOGRAPHY

- Abbeel, P., A. Coates, M. Quigley, and A. Ng. "An Application of Reinforcement Learning to Aerobatic Helicopter Flight." International Conference on Neural Information Processing Systems, Vancouver, British Columbia, December 2006.
- Ahmed, S., N. E. Bajema, S. Bendett, B. A. Chang, R. Creemers, C. C. Demchak, S. W. Denton, J. Ding, S. Hoffman, R. Joseph, E. Kania, J. Kerr, L. Kostopoulos, J. A. Lewis, M. Libicki, H. Lin, K. Miura, R. Morgus, R. E. Odell, E. Pauwels, L. Saalman, J. Snow, L. Steckman, and V. Weber. *Artificial Intelligence, China, Russia, and the Global Order*, edited by Nicholas D. Wright, Maxwell Air Force Base, AL: Air University Press, October 2019.
- Al-Sharman, M., D. Cao, C. LV, and D. Murdoch. "A Sensorless State Estimation for a Safety-Oriented Cyber Physical System in Urban Driving: Deep Learning Approach." *IEEE/CAA Journal of Automatica Sinica*, vol. 8, no. 1, pp. 169–178, January 2021.
- BAE Systems. "BAE Systems Selected by DARPA to Create Autonomy Software for Multi-Domain Mission Planning." <https://www.baesystems.com/en-us/article/bae-systems-selected-by-darpa-to-create-autonomy-software-for-multi-domain-mission-planning>, 3 December 2019.
- Brundage, M., S. Avin, J. Clark, H. Toner, P. Eckersley, B. Garfinkel, A. Dafoe, P. Scharre, T. Zeitzoff, B. Filar, H. Anderson, H. Roff, G. Allen, J. Steinhardt, C. Flynn, S. Ó hÉigeartaigh, S. Beard, H. Belfield, S. Farquhar, C. Lyle, R. Crootof, O. Evans, M. Page, J. Bryson, R. Yampolsky, and D. Amodi. "The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation." arXiv:1802.07228, Cornell University, 20 February 2018.
- Calderon-Arce, C., and R. Solis-Ortega. "Swarm Robotics and Rapidly Exploring Random Graph Algorithms Applied to Environment Exploration and Path Planning." *International Journal of Advanced Computer Science and Applications*, vol. 10, no. 5, pp. 692–702, 2019.
- Carion, N., F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. "End-to-End Object Detection With Transformers." arXiv:2005.12872[cs.CV], Cornell University, 26 May 2020.
- Cox, M. "Meet Ripsaw M5, Textron's Unmanned Battle Wagon." *Military News*, <https://www.military.com/daily-news/2019/10/14/textrons-unmanned-battle-wagon-will-be-able-to-launch-other-robots.html>, 14 October 2019.
- Dafoe, A., E. Hughes, Y. Bachrach, T. Collins, K. R. McKee, J. Z. Leibo, K. Larson, and T. Graepel. "Open Problems in Cooperative AI." arXiv:2012.08630, Cornell University, 15 December 2020.
- Danaher, J. "The Threat of Algocracy: Reality, Resistance, and Accommodation." *Philosophy and Technology*, vol. 29, no. 3, pp. 245–268, 2016.
- Danks, D. "How Adversarial Attacks Could Destabilize Military AI Systems." *IEEE Spectrum*, <https://spectrum.ieee.org/adversarial-attacks-and-ai-systems>, 26 February 2020.
- Dastin, J. "Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women." *Ethics of Data and Analytics*, first edition, chapt. 3, edited by K. Martin, New York: CRC Press, Auerbach Publications, 12 May 2022.
- Defense Advanced Research Projects Agency. "AI Next Campaign." <https://www.darpa.mil/work-with-us/ai-next-campaign>, Arlington, VA, 2018.
- Defense Advanced Research Projects Agency. "DARPA Seeks 'Always on' Interconnected Networks for Multidomain Missions." <https://www.darpa.mil/news-events/2021-05-07>, 7 May 2021.
- Deutsch, D. *The Beginning of Infinity: Explanations That Transform the World*. New York: Viking, Penguin Press, 2011.
- Doncieux, S., N. Bredeche, J.-B., Mouret, and A. E. G. Eiben. "Evolutionary Robotics: What, Why, and Where to." *Frontiers in Robotics and AI*, vol. 2, March 2015.
- Fortune Business Insights. "AI Market Research Report." <https://www.fortunebusinessinsights.com/industry-reports/artificial-intelligence-market-100114>, September 2021.
- Gleave, A., M. Dennis, C. Wild, N. Kant, S. Levine, and S. Russell. "Adversarial Policies: Attacking Deep Reinforcement Learning." ICLR 2020 Conference, Virtual Conference, Formerly Addis Ababa, Ethiopia, March 2020.
- Gonfalonieri, A. "Generative Adversarial Imitation Learning: Advantages and Limits." *Towards Data Science*, <https://towardsdatascience.com/generative-adversarial-imitation-learning-advantages-limits-7c87fc67e42d>, 1 March 2021.
- Grieser, F. "3D Printing Speed: How Fast Can 3D Printers Go." <https://all3dp.com/3d-printing-speed/>, All3DP, 26 January 2016.
- Helfrich, E. "Peraton to Design Autonomous Multi-Domain Network Solution With DARPA." *Military Embedded Systems*, <https://militaryembedded.com/comms/spectrum-management/peraton-to-design-autonomous-multi-domain-network-solution-with-darpa>, 9 February 2022.

BIBLIOGRAPHY, continued

- Hill, S. "AI's Impact on Multilateral Military Cooperation: Experience From NATO." *American Journal of International Law*, vol. 114, pp. 147–151, 27 April 2020.
- Jordan, J. "The Future of Unmanned Combat Aerial Vehicles: An Analysis Using the Three Horizons Framework." *Futures*, vol. 134, December 2021.
- Kamran, D., Y. Ren, and M. Lauer. "High-Level Decisions From a Safe Maneuver Catalog With Reinforcement Learning for Safe and Cooperative Automated Merging." arXiv:2107.07413, Cornell University, 2021.
- Kim, G., and H. Lim. "Reinforcement Learning Based Beamforming Jammer for Unknown Wireless Networks." *IEEE Access*, vol. 8, pp. 210127–210139, November 2020.
- Kissinger, H. A. "How the Enlightenment Ends." *The Atlantic*, <https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/>, June 2018.
- Klare, M. T. "The Challenge of Emerging Technologies." *Arms Control Today*, <https://www.armscontrol.org/act/2018-12/features/challenges-emerging-technologies>, Arms Control Association, December 2018.
- L3HARRIS Technologies, Inc. "The Iver4 900 UUV." <https://www.l3harris.com/all-capabilities/iver4-900-uuv>, 15 March 2022.
- Lang, P., X. Fu, M. Martorella, J. Dong, R. Qin, X. Meng, and M Xie. "A Comprehensive Survey of Machine Learning Applied to Radar Signal Processing." arXiv:2009.13702v1, Cornell University, September 2020.
- Manyika, J., S. Lund, M. Chui, J. Bughin, J. Woetzel, P. Batra, R. Ko, and S. Sanghvi. "Jobs Lost, Jobs Gained: Workforce Transitions in Time of Automation." <https://www.mckinsey.com/~media/mckinsey/industries/public%20and%20social%20sector/our%20insights/what%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/mgi-jobs-lost-jobs-gained-report-december-6-2017.pdf>, McKinsey and Company, McKinsey Global Institute, 6 December 2017.
- Miller, D. "Swarm Warning: The Future of Unmanned Aerial Systems." https://www.army.mil/article/239210/swarm_warning_the_future_of_unmanned_aerial_systems, U.S. Army, 21 September 2020.
- Nashed, S., and S. Zilberstein. "A Survey on Opponent Modeling in Adversarial Domains." *Journal of Artificial Intelligence Research*, vol. 73, pp. 277–327, 2022.
- National Artificial Intelligence Office. "National Artificial Intelligence Advisory Committee (NAIAC)." <https://www.ai.gov/naiac/>, 2020.
- Organisation for Economic Cooperation and Development (OECD). *Artificial Intelligence in Society*. Paris, France: OECD Publishing, 11 June 2019.
- Pane, Y. P., S. P. Nagesh Rao, J. Kober, and R. Babuška. "Reinforcement Learning Based Compensation Methods for Robot Manipulators." *Engineering Application of Artificial Intelligence*, vol. 78, pp. 236–247, February 2019.
- Polyakova, A. "Weapons of the Weak: Russia and AI-Driven Asymmetric Warfare." *Brookings, Series: A Blueprint for the Future of AI*, <https://www.brookings.edu/research/weapons-of-the-weak-russia-and-ai-driven-asymmetric-warfare/>, The Brookings Institution, 15 November 2018.
- PWC. "Annual AI Predictions Survey." <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html>, 2021.
- Qi, D., J. Zhang, X. Liang, Z. Li, J. Zuo, and P. Lei. "Autonomous Reconnaissance and Attack Test of UAV Swarm Based on Mosaic Warfare Thought." The 6th International Conference on Robotics and Automation Engineering (ICRAE), pp. 79–83, Guangzhou, China, November 2021.
- Sculpteo. "Layer Thickness in 3D Printing: An Additive Manufacturing Basic." <https://www.sculpteo.com/en/glossary/layer-thickness-definition/>, accessed 2022.
- Sharre, P., and M. Horowitz. "Artificial Intelligence: What Every Policymaker Needs to Know." <https://www.cnas.org/publications/reports/artificial-intelligence-what-every-policymaker-needs-to-know>, Center for a New American Security (CNAS), 19 June 2018.
- Shea, P. J. "How Is NATO Meeting the Challenge of Cyberspace?" *Prism*, vol. 7, no. 2, pp. 18–29, 21 December 2017.
- Song, Y., and D. Scaramuzza. "Learning High-Level Policies for Model Predictive Control." International Conference on Intelligent Robots and Systems (IROS), Las Vegas, CA, October 2020.
- Starck, N., D. Bierbrauer, and P. Maxwell. "Artificial Intelligence, Real Risks: Understanding—and Mitigating—Vulnerabilities in the Military Use of AI." *Modern War Institute at West Point*, <https://mwi.usma.edu/artificial-intelligence-real-risks-understanding-and-mitigating-vulnerabilities-in-the-military-use-of-ai/>, 18 January 2022.
- Tang, R., Z. Zhimin, Z. Chenlong, and L. Linlin. "The Applications of Artificial Intelligence in Situation Assessment and Game Countermeasure During Unmanned Air Combat." IEEE International Conference on Unmanned Systems (ICUS), pp. 909–913, Beijing, China, October 2019.

BIBLIOGRAPHY, continued

- The AlphaStar Team. "AlphaStar: Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning." DeepMind, <https://deepmind.com/blog/article/AlphaStar-Grandmaster-level-in-StarCraft-II-using-multi-agent-reinforcement-learning>, 30 October 2019.
- Topchyan, A., and T. Topchyan. "Muscle-Based Skeletal Bipedal Locomotion Using Neural Evolution." The 9th International Conference on Computer Science and Information Technologies Revised Selected Papers, pp. 1–6, Yerevan, Armenia, September 2013.
- Trevithick, J. "Boeing Is Building Big Orca Drone Subs for the Navy to Hunt and Lay Mines and More." *The Warzone*, <https://www.thedrive.com/the-war-zone/26513/boeing-is-building-the-navy-big-orca-submarine-drones-to-hunt-and-lay-mines-and-more>, 5 February 2019.
- Trevithick, J. "New Tracked Unmanned Vehicle Bristles With a Whopping 50 Switchblade Suicide Drones." *The Warzone*, <https://www.thedrive.com/the-war-zone/42755/new-tracked-unmanned-vehicle-bristles-with-a-whopping-50-switchblade-suicide-drones>, 15 October 2021.
- Trevithick, J. "This Is Our First Look at How General Atomics' Sparrowhawk Drone Will Get Caught in Midair." *The Warzone*, <https://www.thedrive.com/the-war-zone/41716/this-is-our-first-look-at-how-general-atomics-sparrowhawk-drone-will-get-caught-in-mid-air>, 27 July 2021.
- Tseng, B. "Tesla Shows Us the Risks When Dominant Incumbents Fail to be Bold." *Shield AI*, <https://shield.ai/content/2021/12/6/tesla-shows-us-risks-when-dominant-incumbents-fail-be-bold>, 6 December 2021.
- U.S. Army DEVCOM Army Research Laboratory Public Affairs. "AI Research Helps Soldiers Navigate Complex Situations." https://www.army.mil/article/240917/ai_research_helps_soldiers_navigate_complex_situations, 17 November 2020.
- U.S. Army. "Robotics and Autonomous Systems Strategy." 2017.
- U.S. Department of Defense. *DoD News*. <https://www.defense.gov/News/News-Stories/Article/Article/2640609/memo-outlines-dod-plans-for-ethical-artificial-intelligence/>, 1 June 2021.
- UAS Vision. "Raytheon Gets \$33M U.S. Navy Coyote Contract." <https://www.uasvision.com/2021/03/05/raytheon-gets-33m-us-navy-coyote-contract/>, 5 March 2021.
- United Nations. *UN News: Global Perspective Human Stories*. <https://news.un.org/en/story/2021/11/1106612>, 25 November 2021.
- van Duin, S., and N. Bakhshi. *Perspektiv*. <https://www2.deloitte.com/se/sv/pages/technology/articles/part1-artificial-intelligence-defined.html>, Deloitte, March 2017.
- Vergun, D. "General Says Artificial Intelligence Will Play Important Role in Network Defense." *DoD News*, <https://www.defense.gov/News/News-Stories/Article/Article/2805760/general-says-artificial-intelligence-will-play-important-role-in-network-defense/>, 8 October 2021.
- Walsh, S. M., M. S. Strano, and S. C. Stanton. "Approaching Robotics and Autonomous Systems as an Integrated Materials, Energy, and Control Problem." *Robotic Systems and Autonomous Platforms*, edited by S. M. Walsh and M. S. Strano, pp. xix–xlvi, Cambridgeshire, United Kingdom: Woodhead Publishing, 2019.
- Zhang, B., M. Anderljung, L. Kahn, N. Dreksler, M. C. Horowitz, and A. Dafoe. "Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers." *Journal of Artificial Intelligence Research*, vol. 71, pp. 591–666, 2021.
- Zhang, S., and O. R. Zaiane. "Comparing Deep Reinforcement Learning and Evolutionary Methods in Continuous Control." arXiv:1712.00006v2, Cornell University, 2017.

ARTIFICIAL INTELLIGENCE (AI) FOR WEAPONS SYSTEMS

By Sam Chakour

DSIAC-BCO-2022-216

