Defense Systems Information Analysis Center

AI for Weapons Systems

WEBINAR

SAM CHAKOUR

Outline

- Problem Statement
- AI and Foreign Threats
- Brief History of AI
  - AI and ML
- State-of-the-Art Methods
  - Learning AI
    - Deep Learning
    - Reinforcement Learning
  - Optimization
    - Stochastic Optimization
    - Graph Search Algorithms
  - Emerging AI paradigms
- Application of AI for Weapons Systems
  - Autonomous Systems
  - Perception
  - GNC
  - Mission and Path Planning
  - Intelligent Strategy
  - Opponent Modeling
  - Cognitive Electronic Warfare
- Systems and Programs Applying AI for Weapons Systems
- Summary

Problem Statement

- Intelligence-based warfare, advanced robotics, and hypersonic weapons pose a new breed of qualitative challenges to the U.S. Armed Forces.
- Recent advances in the field of machine learning/artificial intelligence (ML/AI) shed light on the major potential of AI in warfare.
- Cutting-edge offensive and defensive capabilities of AI-enabled weapons:
  - Increase lethality and survivability
  - Improve performance and maintainability
  - Support and automate critical decision making in highly dynamic environments

In its 2021 report, the National Security Commission on Artificial Intelligence urged that the foundations of widespread integration of AI across the DoDmilitary enterprise be in place by 2025.

AI and Foreign Threats

China

- Has the largest share of AI research publications in the world since 2017
- The WIPO ranked China 1st with a 74.7% of the global share of AI-related IP in the past decade
- 12th in 2021 Global Innovation Index
- Actively engaging in the "intelligentization" of its military capability across domains of combat operations

Russia

- Ranks 6^{th} in terms of government strategy
- Advanced from 21^{st} to 16^{th} place in AI research in few years
- In 2019, Russia adopted "National Strategy for the Development of AI" through 2030
- Boosting the private investments in domestic AI industries as a key policy for statewide adoption of AI

AI and Foreign Threats

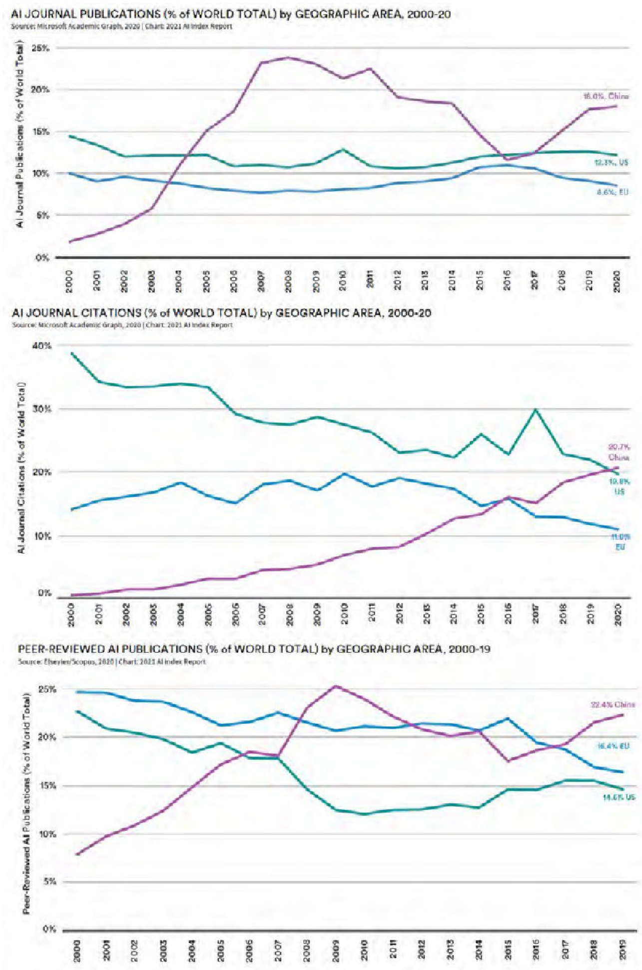

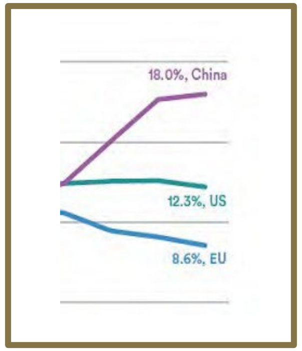

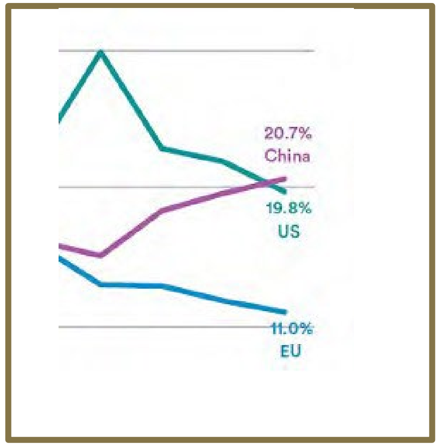

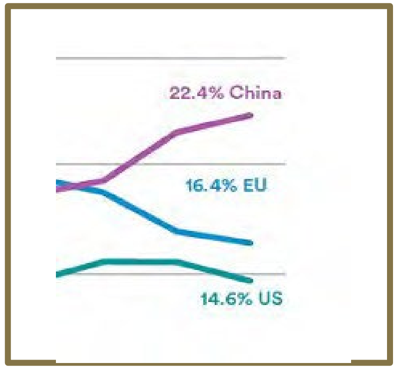

History of AI

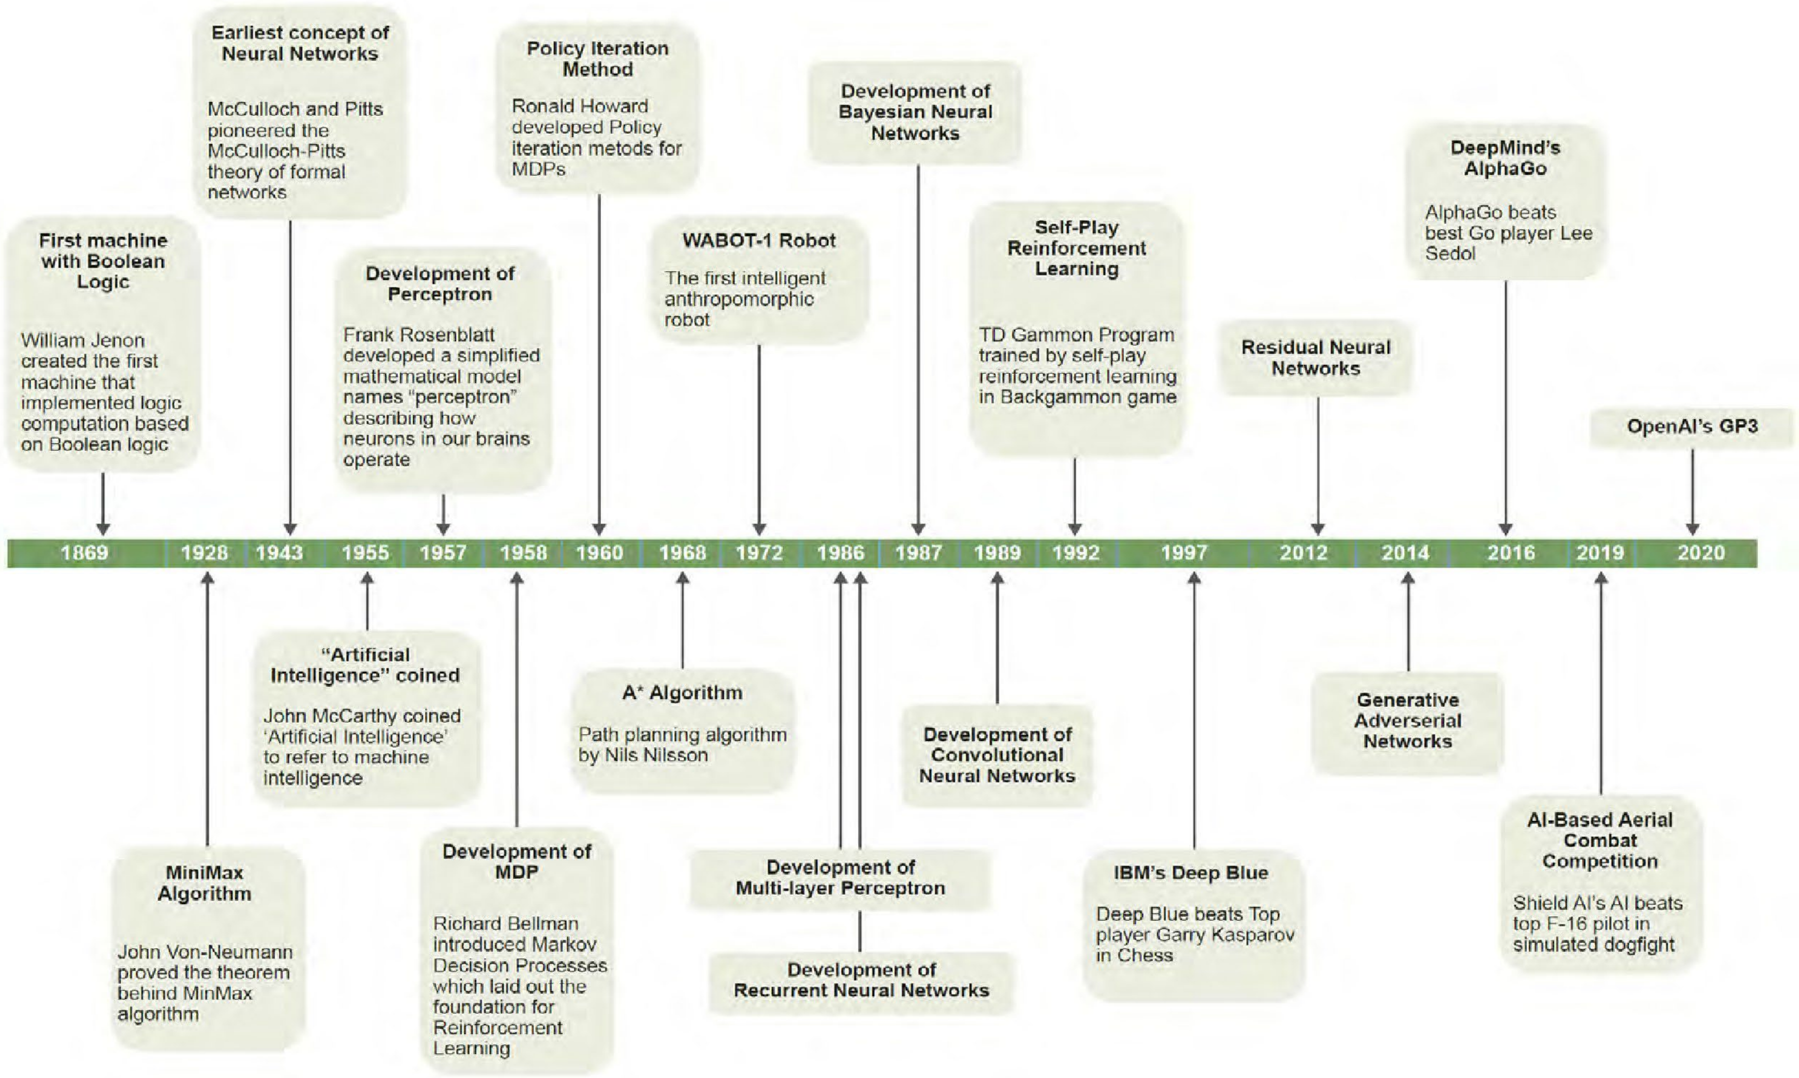

AI Paradigms

The goal of AI is to perceive, reason, construct knowledge, infer, plan, decide, learn, communicate, and efficiently manipulate the environment.

AI, as it is now, is divided into the following paradigms:

- Symbolic AI (e.g., Knowledge-based methods)
- Learning AI (e.g., Machine learning and probabilistic methods)
- Sub-symbolic AI (e.g., Optimization, Search, and Distributed Systems

The focus

Machine Learning (ML) is the ability of a machine to learn from data for the purpose of making accurate predictions.

ML is divided into 4 classes of learning:

1. Supervised Learning
2. Unsupervised Learning
3. Semi-supervised Learning
4. Reinforcement Learning

DSIAC

AI Tiers

Artificial Narrow Intelligence (ANI)

- Context-specific reasoning, need, and inference
- Task-focused autonomous decision making
- Task-focused autonomous planning
- Task-focused autonomous learning
- Task-focused communication

Artificial General Intelligence (AGI)

- Reasoning based on previous experience, need, and inference
- Autonomous decision making
- Autonomous planning
- Autonomous learning and imitation
- Communication in natural language

Artificial Super Intelligence (ASI)

- Reasoning based on previous experience, need, and inference
- Autonomous decision making
- Autonomous planning
- Autonomous learning and imitation
- Communication in natural language

Intelligence

Current State-of-the-Art

State-of-the-Art Methods

State-of-the-Art Methods

Learning AI

1. Deep Learning

Deep learning is the application of deep neural network (DNN) structures to divulge abstract and high-level context from data.

Network Type	Network Architecture	Characteristics	Use/Success
Deep CNN	ResNets	▪ Devised to solve vanishing gradient and degradation problems ▪ Improved performance over previous deep CNN architecture like VGG	▪ Image classification ▪ Object detection ▪ Activity recognition ▪ Machine comprehension ▪ Speech emotion recognition
	DenseNets	▪ Easier to train compared to previous CNN-based architectures ▪ Provides the best representation of images when applied to near-identical images in ImageNet (large visual database)	▪ Image classification ▪ Pattern recognition ▪ Object detection ▪ Image segmentation
	SparseNets	▪ Sparsifies the density to improve performance ▪ Requires a layer depth between 28 and 76	▪ Image classification ▪ Object detection

State-of-the-Art Methods

Learning AI

1. Deep Learning

Network Type	Network Architecture	Characteristics	Use/Success
RNN	LSTMs	▪ Devised to overcome the vanishing gradient problem in classic RNN	▪ Time series processing ▪ Speech recognition ▪ Speech synthesis ▪ Audio processing
	GRUs	▪ Alternative approach to solving the vanishing gradient ▪ Uses less memory and runs faster than LSTMs	▪ Image classification ▪ Scene graph generation ▪ Semantics exploitation ▪ Relation extraction
	NARX	▪ Adds feedback connections to enclose several layers of the network ▪ Adopts an autoregressive model	▪ System identification ▪ Nonlinear filtering
Specialized Networks	Autoencoders	▪ Self-supervised learning network ▪ Unsupervised learning of data representation via encoding, compression, or reconstruction	▪ Dimensionality reduction ▪ Feature clustering ▪ Data compression
	GANs	▪ Generates novel samples from the statistical distribution of the original samples ▪ They can suffer from convergence issues	▪ Interpretation of content from data and generation of novel versions of the data

State-of-the-Art Methods

Learning AI

1. Deep Learning

Network Type	Network Architecture	Characteristics	Use/Success
Specialized Networks	GNNs	▪ Motivated partly by the inability of CNNs to handle non-Euclidean types of data ▪ Based on graph data structure ▪ Inference on graph-represented data having complex relationships and object interdependencies	Applications of structure and unstructured data: ▪ Image classification ▪ Scene graph generation ▪ Semantics exploitation ▪ Relation extraction ▪ Recommender systems ▪ Program reasoning
	DBNs	▪ Based on generative stochastic neural nets that learn probability distribution from input data ▪ Less computationally expensive compared to forward neural nets ▪ Require solid theoretical knowledge of autoencoders and restricted Boltzman machines-based structures	▪ Image classification ▪ Motion capture ▪ Nonlinear dimensionality reductions

State-of-the-Art Methods

Learning AI

2. Reinforcement Learning

Reinforcement learning (RL) models the natural learning process of living systems via action-reward mechanism. RL is formulated for single-agent (SARL) and multi-agent (MARL) settings.

SARL

Actions are behavior an agent can do to change states which are representation of the sensed environment, while rewards are the utility the agents receives for performing desired actions.

The objective is to learn a policy that specifies through time which actions to take from each state to maximize cumulative reward, leading to an overall desired behavior.

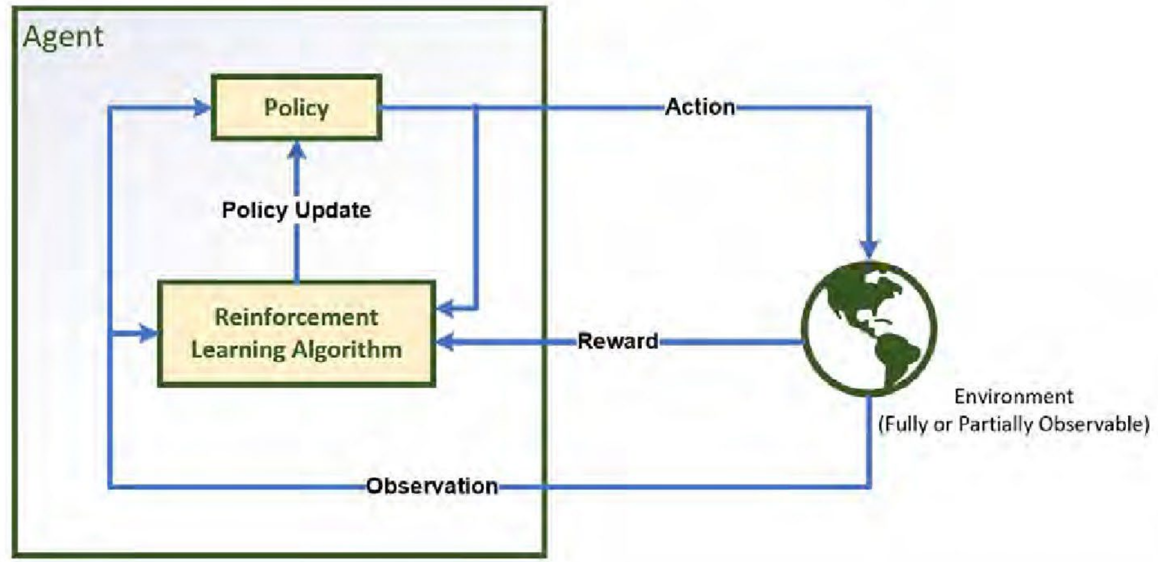

State-of-the-Art Methods

Learning AI

2. Reinforcement Learning

MARL

Concerned with how multiple agents interact with one another and with the environment in which they are all in. They can be formulated to address cooperative, competitive, or adversarial objectives as well as their combination.

Examples

SARL :

- Robotic manipulator learning to organize its environment
- Humanoid learning to walk
- Car learning to self-park

MARL:

- Distributed robotic rescue (cooperative MARL)
- Collaborative manufacturing (cooperative MARL)
- UCAV swarm teaming and targeting (mixed MARL, e.g., cooperative and adversarial)

State-of-the-Art Methods

Learning AI

2. Reinforcement Learning

Quality/Algorithm	SARL	MARL
Advantages	▪ The environment is represented by the uncertainties in dynamics, nonlinearities, disturbances, and model errors. ▪ Can be model free and no prior knowledge of reward function. ▪ Handles stochastic nonlinear dynamics. ▪ Ability to solve for generalized reward function such non-quadratic reward/cost functions.	▪ Experience sharing between agents in the cooperative settings. ▪ Behavior imitation of human or qualified agents. ▪ Inherent redundancy in homogenous MARL which provides increased robustness.

Drawbacks	▪ Requires significant training data and consequently training time.
	▪ For unknown reward function, positive reward is function of the agent's exploration of the environment.
		▪ The combinatorial nature of MARL and multidimensionality of state space pose challenges:
		  ▪ Non-stationarity of the environment due to the concurrent action of other agents.
		  ▪ Non-uniqueness of the learning goal leads to unreachability of the multi-objective equilibrium.
		  ▪ Joint action space increases exponentially leading to scalability.

DSIAC

State-of-the-Art Methods

Learning AI

3. Deep Reinforcement Learning

Deep reinforcement learning incorporates deep network structures with RL algorithms. Deep neural networks interpret inputs and provide predictions of RL policy.

State-of-the-Art Deep RL Algorithms:

Model free:

- Deep Q-learning: e.g., Dueling Deep Q-Network (DQN), Prioritized Dueling DQN
- Deep Policy Gradient Method: e.g., PPO, Actor-Critic, Advantage Actor-Critic, Asynchronous Actor-Critic, and Actor-Critic with Experience Replay
- Deep Deterministic Policy Gradient (DPG)
- Soft Actor-Critic (a bridge between Q-learning and DPG)

Model based:

- Manifest state-based RL
- Latent state-based RL

State-of-the-Art Methods

Optimization and Search Algorithms

1. Stochastic optimization

Stochastic optimization is a family of methods for minimizing/maximizing an objective function when randomness is present in the system.

Heuristics and metaheuristics algorithms are an important categorization of approximate methods of mathematical optimization. They refer to how these algorithms perform approximate optimal solution searches and the types of problems they solve.

- Heuristics are **<u>problem-specific methods</u>** and refer to algorithms that perform informed search to systematically explore search space under a constant heuristic rule.
  - Prone to local optima trap (not optimal or sub-optimal)
- Metaheuristics are <u>not problem-specific methods</u> and refer to high-level strategies for informing and guiding a search using multiple criteria so that search is updated to better explore the search space.
  - Suitable for multi-objective optimization problems

State-of-the-Art Methods

Optimization and Search Algorithms

1. Stochastic optimization

Swarm Intelligence (SI)

The behavior pattern that forms from the decentralized collective actions of distributed, self-organized agents. Each individual agent acts and reacts according to its local rules from which a complex group behavior unfolds. The resultant collective behavior is more advantageous over individual action.

State-of-the-Art Methods

Optimization and Search Algorithms

1. Stochastic optimization

Evolutionary Algorithms (EAs)

Nature-inspired stochastic search algorithms that employ evolution principles – reproduction, genetic crossover, and mutation to improve the outcome of a desired quantity (fitness) in a system.

EAs offer a number of advantages:

- May be applied without expert knowledge of domain-specific heuristics
- Less susceptible to the choice of initial conditions
- Do not suffer reachability issues and, hence, offer global solution (entire solution space is sampled)
- Robust in comparison to gradient-based optimization methods

State-of-the-Art Methods

Optimization and Search Algorithms

2. Stochastic optimization

Physics-Inspired Algorithms

Artificial Potential Field (APF)
A method that models the environment as an artificial potential field and uses virtual force assignment so that target points of interests (e.g., obstacles) generate repulsive fields at these points. APF is used extensively in robotics and aerial path planning and suitable for real-time planning in both static and dynamic environment with unknown and moving obstacles.

Simulated Annealing (SA)
A probabilistic technique in the family of metaheuristic methods for approximating global optimum of a function of arbitrary search space (discrete or continuous), even in cases when the space is replete with local optima. Inspired by the physics of certain materials during heating and cooling phases. SA can solve constrained multi-objective optimization and used for local and global searches.

Quantum-Bacterial Swarm Optimization (QBSO)
A semi-physics 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.

State-of-the-Art Methods

Optimization and Search Algorithms

1. Stochastic optimization

Other Metaheuristics Algorithms

Additional optimization methods within the metaheuristics family of algorithms are those that advocate the 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.

Guided Local Search (GLS)
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.

Tabu Search (TS)
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. A noteworthy property of TS is its applicability to combinatorial optimization, where the objective of obtaining an optimal ordered solution applies.

Reactive Search (ReS)
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.

State-of-the-Art Methods

Optimization and Search Algorithms

2. Graph Search Algorithms

Graph search algorithms are a family of algorithms suitable for path-planning applications, such shortest paths in static and dynamics environment.

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.

RRT algorithm

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. Notable improved variants for faster convergence are RRT*-smart and informed-RRT*.

State-of-the-Art Methods

Optimization and Search Algorithms

2. Graph Search Algorithms

Algorithm/Quality	Type	Advantages	Drawbacks
D*	Informed Search	▪ Suitable for dynamic environment ▪ Real-time replanning ▪ Efficient when dynamic changes happen close to the current node in the search space ▪ More efficient than A*	▪ Slow convergence for higher search space (curse of dimensionality)
RRT	Sample Based	▪ Few heuristics and arbitrary parameters ▪ Provides asymptotic optimality and probabilistic completeness ▪ Faster than A*	▪ Shorter path at the expense of computational efficiency ▪ Slow convergence for higher search space, e.g., 3D path planning ▪ Issue of smoothness of the path

State-of-the-Art Methods

Emerging Paradigms

Neuro-Evolution (NE)

Neuro-symbolic 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.

Advantages

- ❖ Shown to outperform state-of-the-art deep neural networks in video reasoning domains.
- ❖ Improved accuracy.
- ❖ Improved training time.

Drawbacks

- ❖ Coupled complex control flow which makes the computation partly unsuitable for parallelism.

State-of-the-Art Methods

Emerging Paradigms

Neuro-Evolution (NE)

NE is the artificial evolution of neural networks using genetic algorithm (GA). The GA modifies via evolution not only the connection weights but also the network structure leading to an evolved network with enhanced performance.

Advantages

- ❖ Showed significant improvement of models performance
- ❖ A competitive alternative to gradient-based methods for training DNNs and RL models

Application of AI in Weapons Systems

Application of AI in Weapons Systems

Autonomous Systems

Autonomy

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.

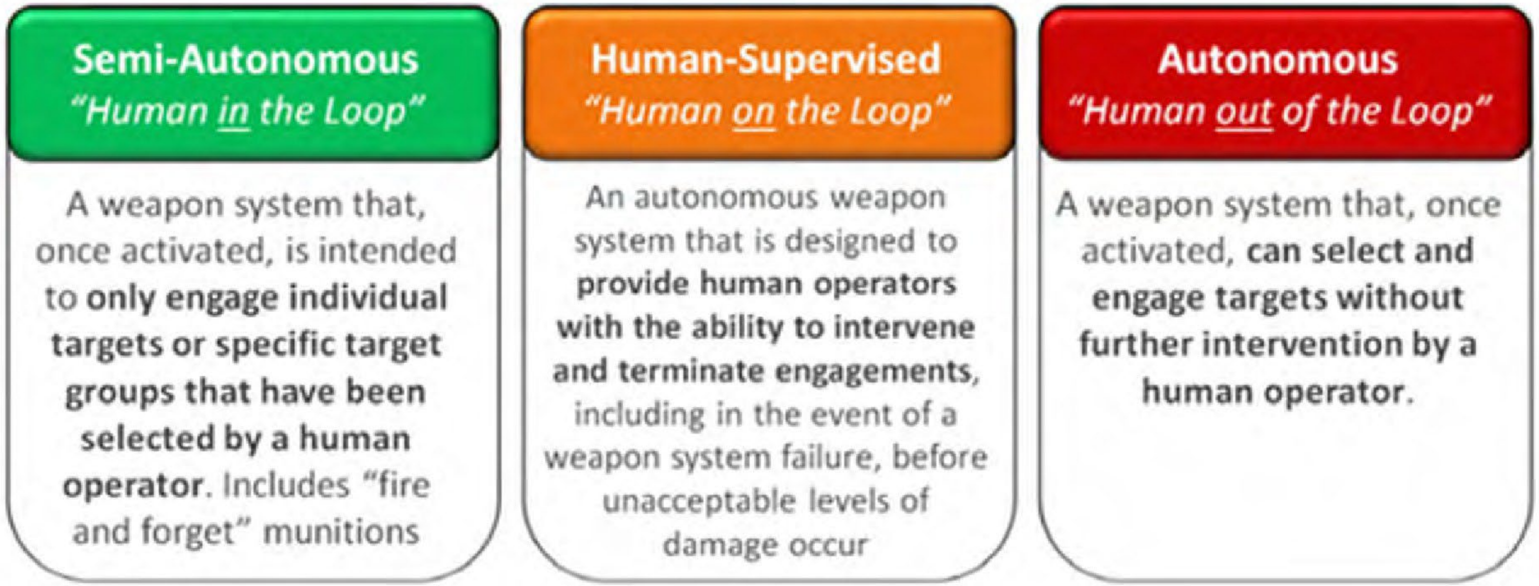

DoD 3000.09 Definitions of Top Levels of Autonomy

Application of AI in Weapons Systems

Autonomous Systems

Autonomy

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.

Levels of autonomy based on the Select, Plan, Act (SP&A) model

Application of AI in Weapons Systems

Autonomous Systems

Autonomy SW stack

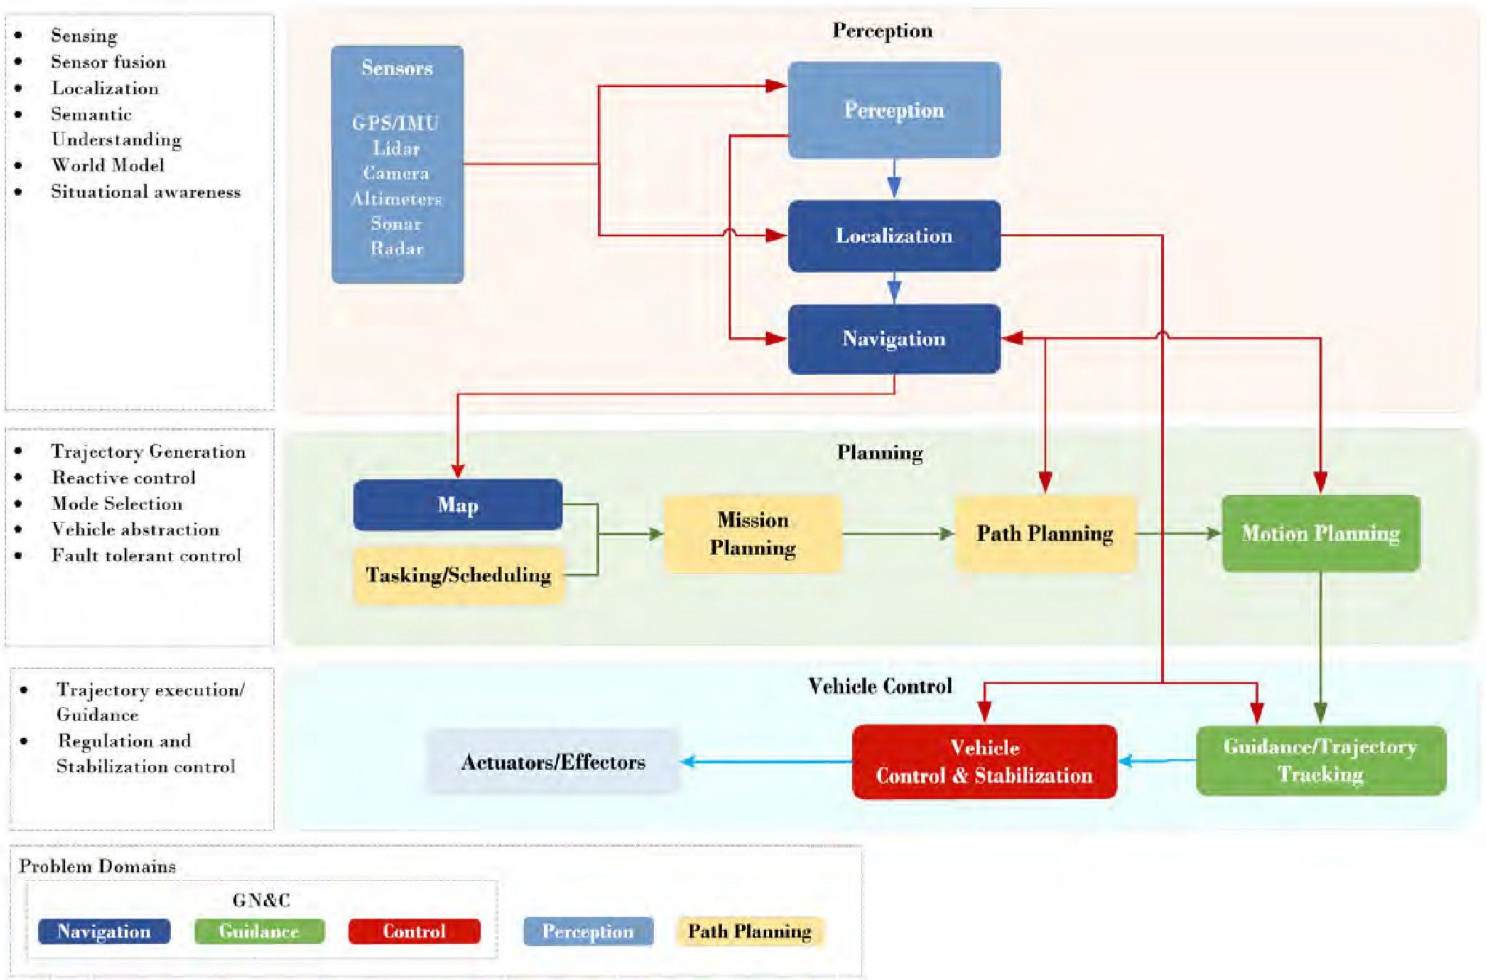

Application of AI in Weapons Systems

Perception

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.

Image Segmentation

Segmentation	Technique	State-of-the-art models
Semantic Segmentation	Object classification + bounding box + Pixel grouping by segmentation mask	DeepLabV2, PSPNet, ParseNet
Instance Segmentation	Semantic segmentation + object detection	R-CNN pipeline (based on ResNet-50)
Panoptic Segmentation	Semantic segmentation + instance segmentation Each pixel is assigned a sematic label and instance ID	End-to-end frameworks: DETR, Detectron2

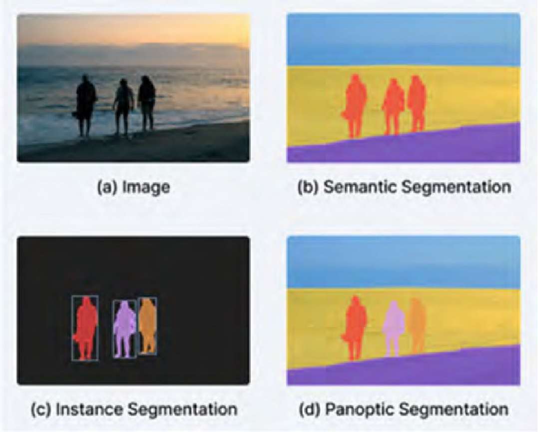

(Source: V7Labs).

- ❖ Semantic Segmentation: Produces high-level segmentation
- ❖ Instance Segmentation: Produces richer output format compared to object detection and sematic segmentation nets separately
- ❖ Panoptic Segmentation: Provides end-to-end object detection

Application of AI in Weapons Systems

Perception

Target Detection, Classification, and Scene Understanding

Image Technique	Description	Application	State-of-the-art
Image Classification	Categorization of what is seen in an image	Identification of pedestrians during car driving, detection of tumors, detection of production defects.	Microsoft AI's Florence-CoSwim-H model (99.02% accuracy)
Image Generation	Generation of artificial images that are indistinguishable from real one	Increase the sample space for training nets for object detection, increase resolution of images, and text-to-image translation	GANs Frechet Inception Distance score of 7.71 was achieved as of 2021 on STL-10 dataset
Pose Estimation (Humanoids)	Estimation of human body posture and joint kinematics	Activity analytics, crowd monitoring, target surveillance.	Latest deep nets (e.g, ResNet-50) achieved 99.5% human pose estimation accuracy on the Leeds Sports Poses dataset
Image Segmentation	Assignment of individual image pixels a class or category for classification	Self-driving vehicles (e.g., UGV), image analysis, scene understanding	Top performing DNNs topologies (e.g., R-CNN) achieved 86.20% accuracy on challenging Cityscapes dataset
Visual Reasoning	Inference from a combination of visual and textual data	Abstract reasoning, intent inference about action in image	Ongoing and challenging area of research
Visual Q&A (VQA)	Combines language understanding, vision, and common-sense reasoning to answer open-ended question about a scene at a high-level	Explainable AI	Top performing DNNs topologies achieved 79.8% in 2021 VQA challenge close to human's 80.8% baseline. 89.6% accuracy was achieved in Kinetics-600 dataset (Google and Brown University)

Application of AI in Weapons Systems

Perception

Target Detection, Classification, and Scene Understanding

Video Technique	Description	Application	State-of-the-art
Object Detection	The ability to identify objects within an image Classification + localization + masking	Target recognition, target tracking	More accurate: Faster R-CNN, Cascade R-CNN, Mask R-CNN Faster: YOLO, RetinaNet, SSD YOLO : 80.7% mean average precision (MAP) (in 2021) Faster R-CNN : 87.69 MAP (in 2021)
Activity Recognition	The ability to identify high-level abstract activities in video	Recognition of activities (good or bad)	Dynamic GCN, CNN-LSTM Top-1 (2022) achieved 89.6% accuracy on Kinetics-600 dataset, 89.10% on kinetics-400 dataset, and 82.25 on Kinetics-700 dataset
Visual Commonsense Reasoning (VCR)	Aims at answering questions about a scenario presented from image frames and providing reasoning behind the answers	Explainable AI, intent prediction (see opponent modeling and intelligent strategy).	The 2021 best mark on VCR is about 72%, still below human baseline of 85%. Improvements have become increasingly marginal, which suggests new techniques may be required to further improve performance closer to human baseline.

Application of AI in Weapons Systems

Guidance, Navigation, and Control

Intelligent Control

Intelligent control refers to hybrid control techniques (AI-enabled) combining ML with control theoretic-based techniques for the control of complex, nonlinear, dynamic systems. The goal of intelligent control is to provide an alternative approach for classic adaptive nonlinear control of complex dynamics.

	Technique	Description	Applications
Intelligent Control	DNNs	Used as multivariable mapping/function approximator. When combined with adaptive control or nonlinear control, hybrid methods provide improved online adaptation to time-varying dynamics	▪ Nonlinear control of landing in unsteady aerodynamics ▪ Re-entry maneuvering and guidance ▪ Accurate robotic trajectory control
	Reinforcement Learning	Suitable for the action control of stochastic processes. It is applied to higher level motion control (outer loop control).	▪ Guidance of hypersonic re-entry vehicles ▪ Improved guidance of UAV toward a target ▪ Adaptive control of quadruped robots
	Imitation Learning	Offers an optimal policy by imitating the expert decision or actions. Imitation learning is useful when an expert's desired behavior demonstration is easier than specifying a reward function (unobtainable or sparce reward function).	▪ Bipedal robotic control from visual human walking ▪ Autonomous flight control
	Deep Model Predictive Control	Deep model prediction control (DeepMPC) is an approach to model learning for predictive control designed to handle variations in a system's environment and variations during system actions	▪ Trajectory tracking of unmanned vehicles ▪ Control of high-dimensional systems

Application of AI in Weapons Systems

Guidance, Navigation, and Control

Intelligent Control

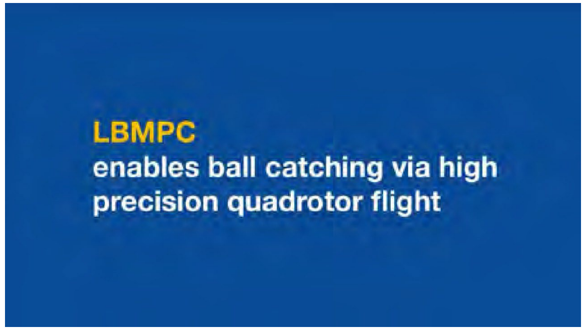

Deep Learning + MPC

Application of AI in Weapons Systems

Guidance, Navigation, and Control

Localization and Navigation

Problem domain	State-of-the-art techniques	Applications
Localization	ORB-SLAM, Dense Piecewise Planar Tracking and Mapping, SONAR-SLAM	▪ Autonomous systems localization in GPS-denied environment ▪ Robotic navigation in urban or indoor environment
Navigation	Unscented Kalman Filtering (UKF), Square-root-cubature Kalman Filtering (SCKF), Particle Filtering (PF)	▪ Autonomous systems navigation ▪ Information/sensor fusion

System Identification

Problem domain	State-of-the-art techniques	Applications
System Identification	NARX, LSTM	▪ Modeling of dynamic systems ▪ Timeseries forecasting ▪ Filtering

Application of AI in Weapons Systems

Mission and Path Planning

Path planning problem computes the optimal path or quasi-optimal 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.

Path planning/search methods are solved by:

- Heuristics and metaheuristic algorithms
  - Swarm intelligence
  - Evolutionary methods
- Intelligent schemes, e.g., reinforcement learning

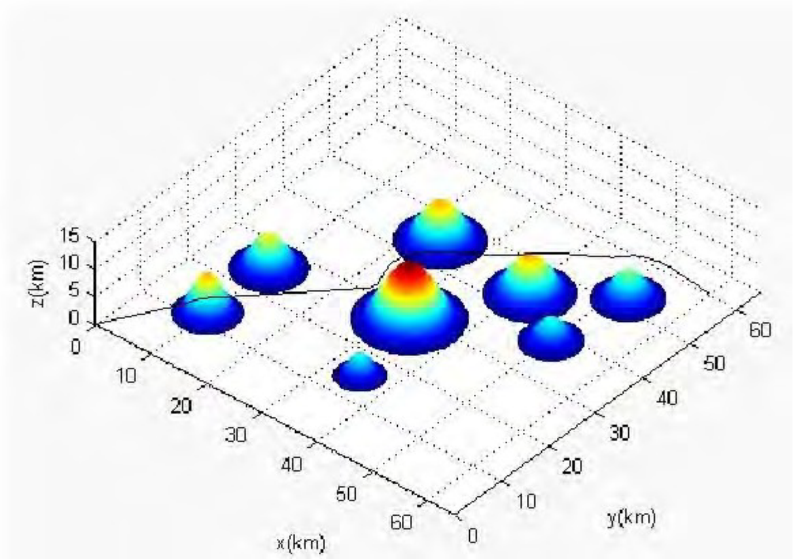

Artificial Potential Field algorithm-based UAV path planning in dynamic environment

Application of AI in Weapons Systems

Mission and Path Planning

Technique	Algorithm	Characteristics	Applications Examples
Physics-Based Metaheuristics	Artificial Potential Field (APF) and variants	▪ Good convergence property	▪ Path planning ▪ Online re-routing
Evolutionary Algorithms	Genetic Algorithm, Memetic Algorithm, Evolutionary Strategy	▪ Global and good robustness property ▪ Not suitable for fine search	▪ Path planning and decentralized control of a swarm of UAVs
Swarm Intelligence	Particle Swarm Optimization	▪ Good convergence but susceptible to premature convergence towards local optimum	▪ Path planning ▪ Parameter tuning of control algorithms ▪ Load balancing
	Ant Colony Optimization	▪ Suitable for discrete optimization problems ▪ Computationally parallelizable ▪ Good robustness even against noisy sample space	▪ In-flight replanning of UAVs and obstacle avoidance
	Multi-Objective ACO	▪ Good robustness even against noisy sample space	▪ Task/Mission allocation of heterogenous UAVs with multiple objectives (task benefit, UAV damage, total range) under physical and operational constraints
	Wolf-Pack Algorithm	▪ Good convergence property ▪ Good robustness property ▪ Good accuracy and stability	▪ Multi-target search problem in an unknown environment

Application of AI in Weapons Systems

Mission and Path Planning

Technique	Algorithm	Characteristics	Applications Examples
Swarm Intelligence	Artificial Shepherding	▪ Suitable for human-swarm interactions ▪ Only the shepherd is trained instead of the whole swarm with the control behavior ▪ Combines rule-based and learning-based algorithms, which allows for adaptation ▪ Suffers from scalability issue of the swarm	▪ Used in conjunction with RL for swarm behavior control
	Spider Monkey Optimization	▪ Super overall performance ▪ Strong robustness ▪ Stable solutions ▪ Faster convergence rate	▪ UAV path planning ▪ Path re-routing
	Fruit-Fly Algorithm	▪ Susceptible to premature convergence ▪ Available variants address local minima convergence problem	▪ UAV path planning in environment with multiple threat sites and constraints
	Chicken Swarm Optimization	▪ Good convergence speed and accuracy ▪ Good global optimization performance ▪ Susceptible to premature convergence ▪ Available variants improve global convergence	▪ Parameter estimation of nonlinear system

Application of AI in Weapons Systems

- Intelligent Strategy

System	Bot/Algorithm	Description	Techniques	Property
Single-agent	AlphaGo (Deepmind)	Plays Go strategy board game	▪ CNN for perception and RL-based for decision-making. ▪ Monte-Carlo Tree Search (MCTS) governed by policy and value networks	▪ Trained from 30 million expert move
	AlphaGo Zero (Deepmind)	Plays Go strategy board game	▪ Uses ResNet to predict policy and values from given state ▪ ResNet evaluates the best move instead of MCTS rollout	▪ Trained by RL from self-play using only board as inputs and game rules. This means it is possible to train an agent from zero capability to expert in complex setting using only rules.
	AlphaZero (Deepmind)	Plays Go strategy board game	▪ Similar to AlphaGo Zero	▪ Generalized to Chess and Shogi games
	MuZero (Deepmind)	Plays Go strategy board game	▪ Uses ResNet and RL	▪ No prior knowledge of the rules or environment dynamics - completely tabula rasa ▪ It learns an accurate model of an environment dynamics which it uses to plan the best actions
Multi-agent	Five (OpenAI)	Plays DOTA2 (Multiplayer battle video game)	▪ Policy-gradient Multi-agent RL	▪ Trained by RL from self-play (180 years of self-play) ▪ Allows for competitive and collaborative actions in an imperfect information, zero-sum game environment ▪ Intent prediction and anticipatory actions
	AlphaStar (Deepmind)	Plays StarCraft II (Multiplayer battle video game)	▪ Policy-gradient based multi-agent RL ▪ Minimap features are extracted by RNN ▪ Time sequence of observations is processed by an LSTM	▪ Initial training by supervised learning from thousands of human replays. ▪ Allows for competitive and collaborative actions ▪ Intent prediction and anticipatory actions

Application of AI in Weapons Systems

Opponent Modeling

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.

Setting	Technique (Examples)
Two-player fighting game	Combined techniques: ▪ Opponent modeling via Q-learning ▪ Best response selection via Enhanced Rolling Horizon Evolution Algorithm (E-RHEA)
Multi-player strategy battle game	Combined architecture: ▪ Convolutional LSTM to explore spatio-sequential correlations in large training data from the game (Brood War game)

Soccer game	▪ Bayesian Policy Reuse (BPR) for nonstationary opponents in Markov games:
	    ▪ Off-line learning of opponent policy
	    ▪ Belief function used to measure similarity between each approximated opponent policy

Application of AI in Weapons Systems

Opponent Modeling

Setting	Technique (Examples)
Nonstationary zero-sum Markov games (opponent strategy changes concurrently)	▪ Opponent modeling stage and heuristics integrating eXtended Classifier System (XCS) for action selection and policy learning
Misdirection and counter-misdirection in nonstationary competitive setting	▪ Granovetter threshold model (flock behavior in response to group communication)
Nonstationary Markov games	▪ Opponent modeling via DNN approximation ▪ Opponent policy is determined using an Extended Deep Bayesian Policy Reuse algorithm (Deep BPR+)
Play calling in football	▪ Reinforcement learning to find utilities and optimal policy from value iteration and greedy policy computation at each state.

Application of AI in Weapons Systems

Cognitive Electronic Warfare

Cognitive electronic warfare is the field of application of cognitive techniques and AI in electronic warfare systems.

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

AI technique apply, in particular, to:

- Automatic modulation classification
- Automatic intra-pulse modulation classification
- Radar pulse repetition interval tracking

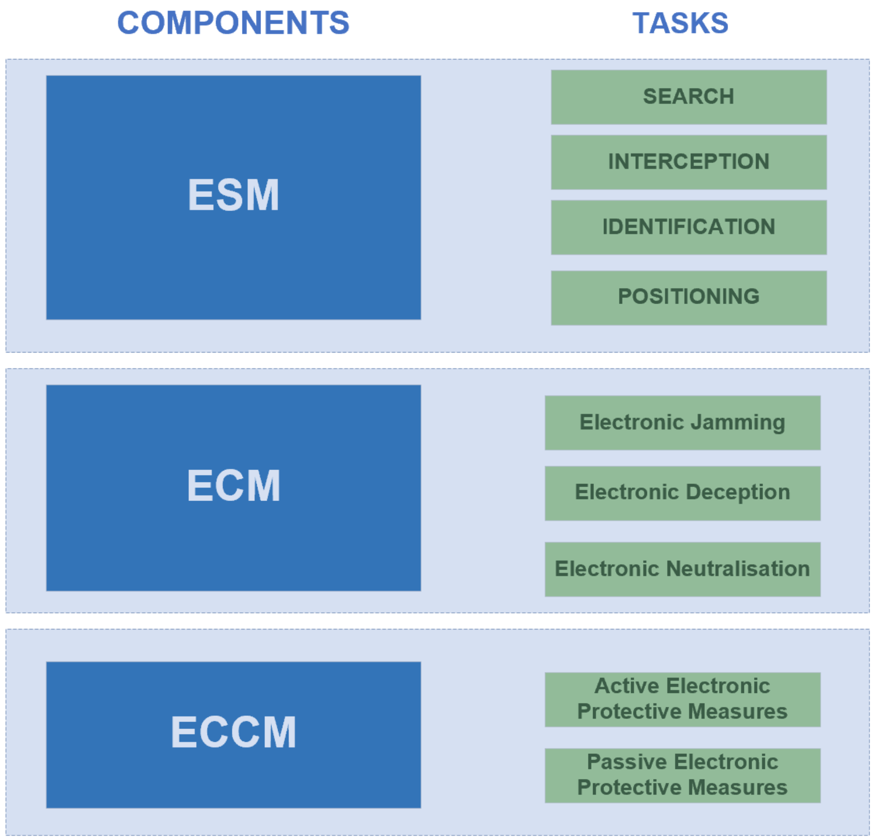

EW top-level functional system components

Application of AI in Weapons Systems

Cognitive Electronic Warfare

Component	Techniques	Application
ESM	ResNets	▪ Automatic classification of 24 types modulation signals
	Ward Clustering and Probabilistic Neural Nets	▪ Classification of received radar pulse signals
	LeNet-5	▪ Automatic intra-pulse modulation classification (96.1% detection rate)
	R-CNN (tiny-YOLO)	▪ Signal detection from spectrogram (trained from 20k event containing spectrograms)
	LSTM	▪ Detection of radar-jamming signals (> 92% detection accuracy)
ECM	SVM	▪ Jamming-style selection
	LSTM	▪ Deception by detection and tracking of pulse repetition interval (RPI) of enemy radar signal ▪ Saturation of spectrum with fake returns to raise enemy's radar PFA rate

Application of AI in Weapons Systems

Cognitive Electronic Warfare

Component	Techniques	Application
ECM	RL (Multi-armed Bandit-based learning)	▪ Optimal jamming strategy without a-priori knowledge of opponent strategy or channel gains
ECCM	Deep RL	▪ Intelligent jamming against enemy's cognitive radio
	DQN	▪ Frequency hopping for a cognitive radar without a-priori knowledge of jamming model (online learning of enemy's jamming strategy)

DSIAC

AI-Enabled Weapons Programs and Systems

AI-Enabled Weapons Programs and

Systems AI-Enabled Weapons Programs

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.

U.S. Air Force's Vanguard Programs
The U.S. Air Force's program objective is to demonstrate networked collaborative and autonomous (NCA) weapons by creating an integrated weapon system of systems where weapons autonomously work together to increase survivability and lethality.

U.S. Air Force's Skyborg Program
As one of Vanguard programs, Skyborg 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. Skyborg will fast-track the capabilities toward MUM-T and resilient autonomy, thereby improving survivability and lethality against near-peer adversaries.

AI-Enabled Weapons Programs and Systems

AI-Enabled Weapons Programs

The U.S. Army's Constructive Machine-learning Battle with Adversary Tactics (COMBAT) program The 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 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.

DARPA's Adapting Cross-Domain Kill-Web (ACK) program
The 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.

AI-Enabled Weapons Programs and Systems

AI-Enabled Weapons Programs

The U.S. Air Force's Kaiju program
The program seeks to create an EW ecosystem to demonstrate next generation of systems capable of autonomous electronic attacks and enemy air defenses countermeasures.

DARPA's Angler Program
The 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 semi-opaque surroundings.

The Office of the Secretary of Defense's Maven Project
The program 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.

AI-Enabled Weapons Programs and Systems

Aerial Systems

- Shield AI Hivemind
  - An AI pilot capable of close aerial combat, teaming, and GPS denied environment
- Shield AI V-Bat
  - VTOL UCAS capable of various mission operation such as infantry clearance, air-defense breach, and swarm operation thanks to Hivemind autonomy core
- Kratos XQ-58 Valkyrie
  - A technology demonstrator UCAS for Skyborg's open autonomy SW architecture. On March 2021, the demonstrator successfully launch Altius-600 UAS
- MQ-20 Avenger
  - Skyborg team demonstrated the portability and scalability of Skyborg autonomy software on MQ-20 as another technology platform for heterogenous teaming
- Autonomous Loitering Munitions
  - Lightweight, multidomain capable, and ground launched autonomous munition designed to autonomously loiter and prosecute acquired targets
- Dynetics X-61 Gremlins
  - Technology demonstrator designed to provide semi-autonomous operation, mid-air recoverable, and low-cost UCAS solution

(Source: Shield AI Used With Permission, https://shield.ai/products)

(Source: Shield AI Used With Permission, https://shield.ai/products)

DSIAC

AI-Enabled Weapons Programs and Systems

Maritime Systems

- Sea Hunter
  - Sea Hunter is an unmanned surface 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.
- Seahawk
  - A newer and scalable version of the Sea Hunter that was launched in August, 2020

(Source: U.S. Navy)

AI-Enabled Weapons Programs and Systems

Land Systems

- QinetiQ/Pratt Miller's EMAV
  - Expeditionary Autonomous Modular Vehicle (EMAV) is a large robotic combat vehicle (RCV-L) designed to meet US Army requirements for the RCV decisive lethality continuum concept. The system is being verified for its support in MUM-T operation environment.
- Textron System's 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.
- Rheinmetall's Lynx KF41
  - KF41 is an optionally piloted infantry fighting vehicle under development that integrates 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.

QinetiQ/PM EMAV V-L (Source: QinetiQ [159]

Ripsaw M5 (Source: U.S. Army Photo Courtesy of Textron Systems,

AI-Enabled Weapons Programs and Systems

Swarm Systems

❑ 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.

❑ Mako UTAP22

- Developed to specifically operate in tactical MUM-T and carry out swarm operations. It embeds Skyborg's autonomy core and can coordinate commands for attack and maneuvering from ground-based and air-based C2.

❑ Collaborative Small-Diameter Bomb Swam (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.

❑ Riptide Swarm Micro-UUVs

- Riptide micro of UUVs are highly flexible platforms that integrate maritime open architecture, capable of carrying autonomous sorites, maintain underwater communication, and exchange acoustic target intelligence data to engage surface targets.

❑ Coyote UAS Block 3

- Coyote is a tube-launched, small autonomous, expandable, and multi-mission UAS adaptable for diverse individual or swarm missions, including EW, surveillance, and strike.

❑ CARACaS Swarm

- CARCaS is an autonomy core framework for multiagent coordination and control that has been used by the Navy to create "swarmboats." The system provides fully distributed, multiagent operations. The cooperative behavior models supported by the autonomy core are patrol, track, inspect, and trail.

(Source: NASA).

(Source: Navy)

AI-Enabled Weapons Programs and Systems

ISR and Targeting Systems

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) addressing the requirement for processing, exploitation, and dissemination (PED) of data. 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.

MQ-9 Reaper (Source: Photo by Airman 1st Class William Rosado)

Nemesis

HRL Labs' NEuroMorphic EyeS In the Sky (NEMESIS) is a bioinspired visual AI system 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.

Nemesis (Source: Getty Images

Summary

Swarm Intelligence, human-robot interaction, planning and control, 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.

Questions

References

1. DSIAC State-of-the-Art Report: "Artificial Intelligence for Weapons Systems," September 2022