ARTIFICIAL INTELLIGENCE
for land operations

GIGAT

ARTIFICIAL INTELLIGENCE

FOR LAND OPERATIONS

Contributors
Julien Bzowski (Safran.AI)
Maud Foresti (Arquus)
Christophe Guettier (Safran)
Arnault Ioualalen (Numalis)
Eric Lebigot (Bertin Technologies)
Guillaume Quin (MBDA)
Bruno Ricaud (KNDS FR)

Chairs
Juliette Mattioli (Thales)
Michel Bouvet (GICAT)

ARTIFICIAL INTELLIGENCE FOR LAND OPERATIONS

SUMMARY

Developed from the work of a GICAT working group, this report is a collective advocacy document in favor of expanding the use of artificial intelligence (AI) for land operations. It addresses operational use cases, available methods, and the responses that the French defense industry must provide in order to improve force effectiveness and ensure that the French Defense Industrial and Technology Base (BITD) remains competitive.

On the one hand, the rapid evolution over the past several years of optimization techniques, multiprocessing capabilities, data storage capacity, and machine learning methods has significantly accelerated the development of artificial intelligence. This report explains the different approaches, methods, and techniques of AI in the most accessible way possible, while seeking to "demystify" AI.

On the other hand, the military context has changed profoundly. Modern conflicts are characterized by the increasing density of threats, the proliferation of low-cost drones, the acceleration of OODA loops, and the need to achieve simultaneous information, decision, and operational superiority.

AI is no longer a prospective option—its operational deployment has become critical.

Drawing on a set of specific capability examples, the report illustrates possible uses of AI, its added value, and the remaining efforts required to bring AI to its full potential and further improve operational effectiveness. It identifies eleven major use cases: radar detection of aerial threats, electro-optical/infrared classification, the establishment of a tactical picture from ISR drones, decision support in classified environments, reduction of operator cognitive workload, autonomous vehicles, maneuver support, predictive maintenance, military logistics, simulation-based training, and detection of loitering munitions. These applications draw on the full spectrum of AI paradigms—connectionist, symbolic, and generative—depending on the maturity level and domain-specific constraints of each field.

Three structuring responses are proposed to address the limitations of current approaches.

The first is **hybridization**. Faced with the opacity of purely connectionist systems and the fragility of purely symbolic systems, the most promising path lies in combining the two. Neuro-symbolic approaches, Physics-Informed Neural Networks (PINNs), which constrain learning through physical laws, and Geometry-Informed Neural Networks (GINNs) illustrate this convergence. Hybrid AI simultaneously offers robustness, explainability, data frugality, and the ability to comply with doctrinal or regulatory constraints—properties that are indispensable in a critical defense context.

The second response concerns **trust and performance assurance**. After recalling the six requirements of the European AI Act—robustness, effectiveness, reliability, usability, transparency, and human oversight—this report emphasizes that these requirements take on a particularly demanding form for defense systems. Safety may require formal validation or even certification; real-time explainability becomes a condition for operational acceptability; and cybersecurity maintains a complex bilateral relationship with AI, as AI is both vulnerable to adversarial attacks and a lever for anomaly detection. It is not enough for a model to be accurate: it must be demonstrably robust, consistent, and controllable.

The third response concerns **integration and human adoption**. The integration of AI functions into embedded systems requires trade-offs in compression, latency, and SWaP—size, weight, and power. Qualification must address the actual behavior of the system in its deployment environment, not merely that of a laboratory prototype. In addition, operator adoption determines the real effectiveness of these systems. Explainability is not a luxury: it is the condition for a trusted human-machine relationship that keeps the human at the center of final decision-making, particularly when lethal force is involved.

In summary, this report calls for a systemic and sovereign engineering discipline for trustworthy AI, combining paradigm hybridization, incremental qualification, and the federation of BITD stakeholders around common methodologies.

1. Introduction

The very rapid evolution in recent years of optimization techniques, multiprocessor computing capabilities, data-storage capacities, and machine-learning methods has significantly fostered the development of artificial intelligence (AI). It is worth recalling the original definition of AI: "*a set of theories and techniques enabling an artificial system to simulate intelligence.*" The properties of intelligence encompass a broad range of cognitive capabilities used to manipulate data, information, and knowledge, including perception, learning, reasoning, decision-making, action, and knowledge.

AI (Artificial Intelligence)
A set of techniques enabling a machine to perform tasks that normally require human intelligence: perceiving, reasoning, deciding, and learning. AI is not a single system but a diverse family of methods, including symbolic, statistical, and generative AI. Example: software that automatically analyzes satellite images to detect vehicles or troop movements.

1.1. Criticality of the Need

Land operations are now characterized by several categories of requirements: requirements for data, information, and knowledge superiority, enabling better action planning and greater transparency of the battlespace; requirements for denying enemy access across all multi-domain and multi-field components; and requirements for decision and operational superiority, such as speed of action, multi-entity and multi-domain collaboration, massification of forces and effects, and increased lethality of actions. Lessons learned from recent and ongoing conflicts clearly show the extent to which systems augmented by AI-based capabilities help meet these requirements. The point has been reached where delays in implementing these capabilities can substantially alter the balance of power and change the outcome of a battle. AI-based capabilities have therefore become critical to achieving information, decision, and operational superiority in the air-land context and in other domains.

Their contribution is distinctive in each of the above dimensions:

- **Transparency**: AI solutions are more effective than limited numbers of human operators at processing the constantly growing mass of heterogeneous data generated by the battlefield. AI-based tools simplify the representation and cross-correlation of such data and the extraction of relevant information. They thereby make it possible to establish and understand, faster and holistically–in the sense of considering the object as a whole–a shared tactical or operational situation. This helps lift the fog of war, facilitates situation assessment, and enables the anticipation and execution of effective defensive or offensive maneuvers under time constraints.
- **Denial**: AI-enhanced systems give individual platforms a capacity for continuous action independent of human physiological limits, including when no human can operate them any longer. AI is thus foundational to a new capability for intelligence or strike vectors, allowing them to navigate and carry out missions under jamming conditions affecting GNSS and communications–conditions that are becoming generalized both at the line of contact and in the adversary's depth.
- **Speed**: AI accelerates battlefield tempo in two ways. First, it contributes to the automation of functions that process information and trigger actions–such as self-protection–at reflex speed. Second, because of its inherent ability to process massive information flows, it accelerates the OODA loop: faster perception of the environment, automatic dissemination of relevant information to the appropriate echelons, rapid simulation and assessment of response scenarios, and activation of relevant effectors under critical time constraints.1
- **Collaboration**: AI contributes in two ways. First, AI-enabled processing simplifies the representation of the environment and therefore makes it possible to share lighter vectorized information in contexts of denial or reduced connectivity. Second, software systems integrating AI enable the autonomization of combat platforms, allowing machine-to-machine collaboration and opening the way to new effects, such as multi-sensor surveillance systems or cooperative area-defense systems.
- **Mass:** AI enables the multiplication of sensor-data processing and synchronized effector actions. Without AI-augmented capabilities, the use of waves or swarms of drones is limited by the human ability to manage only a small number of vectors simultaneously. By contrast, automation and greater autonomy enabled by AI-integrated systems allow a single operator to manage multiple vectors while adapting to complex, non-

[1] The OODA loop–"Observe, Orient, Decide, Act"–is a conceptual decision-making model developed by Colonel John Boyd, a military strategist and fighter pilot in the United States Air Force. It is designed to improve agility and effectiveness in dynamic and uncertain environments.

preprogrammed situations. The human, material, and financial cost of operations involving massed effectors or sustained presence in an area is thereby reduced, expanding the range of options available to command.

- **Lethality:** AI augmentation of combat assets makes it possible both to deploy more attributable vectors–to wear down the enemy and erode combat potential–and to make these vectors more effective through faster, more precise, and/or more systematic targeting. This enables increased lethality of offensive means, such as AI-assisted targeting systems, or enhanced protection in a defensive framework, such as counter-UAS systems. Behind this question lies the issue of trust and the reliability of targeting.
- Finally, AI improves the **effectiveness of military strategy** across three complementary levels – strategic, operational, and tactical – by accelerating planning, optimizing the relevance of monitoring, and improving the quality of tactical operations within a synchronized framework.

Having fundamental AI skills was already essential five years ago; being able to deploy AI operationally is now critical. This report will therefore seek to: (1) set out the principles governing the design of AI solutions; (2) explain their main uses and identify priorities; (3) present the challenges currently facing their implementation; and (4) propose solutions.

1.2. Definitions and Vocabulary

The design of AI solutions relies extensively on the exploitation of data, information, and knowledge. These terms must therefore be defined first. In the remainder of this document, the concepts of "data," "information," and "knowledge" are used as follows.

Data is a raw element that has not yet been interpreted or contextualized. It is the direct result of a measurement that may be collected by a tool or by a person, or may already be present in a database. Data therefore includes numerical, symbolic, textual, or logical data, as well as software components such as code and executables.

Structured Data
Structured data is organized according to a predefined and rigid format, and can be stored in tables where each item of information occupies an identified column. It is directly usable by conventional IT tools. Important classes of AI algorithms require such data as input, such as random forests. Examples include a personnel-management table with fixed columns–rank, unit, assignment date, specialty–or a maintenance history listing, for each vehicle, the date, type of intervention, and technician responsible.
Unstructured Data
Unstructured data does not follow a simple predefined format and cannot conveniently be stored in a table. It now represents the bulk of the world's data volume, often estimated at more than 80 percent. Its exploitation requires specific AI techniques, such as natural-language processing or computer vision. Examples include operational reports written in free text, drone video feeds, audio recordings of communications, aerial photographs, and messages on social networks monitored as part of adversary-watch activities.

Information is intelligible data that takes on meaning through structuring. Information is therefore, by definition, pre-interpreted semantic data, usually by a programmer. In other words, contextualizing data creates added value and turns it into information.

Knowledge is the result of reflection on analyzed information. Unlike information, knowledge can be shared and rests on a collective reference framework, such as a business or domain semantics. It is also possible to perform different types of logical reasoning over knowledge.

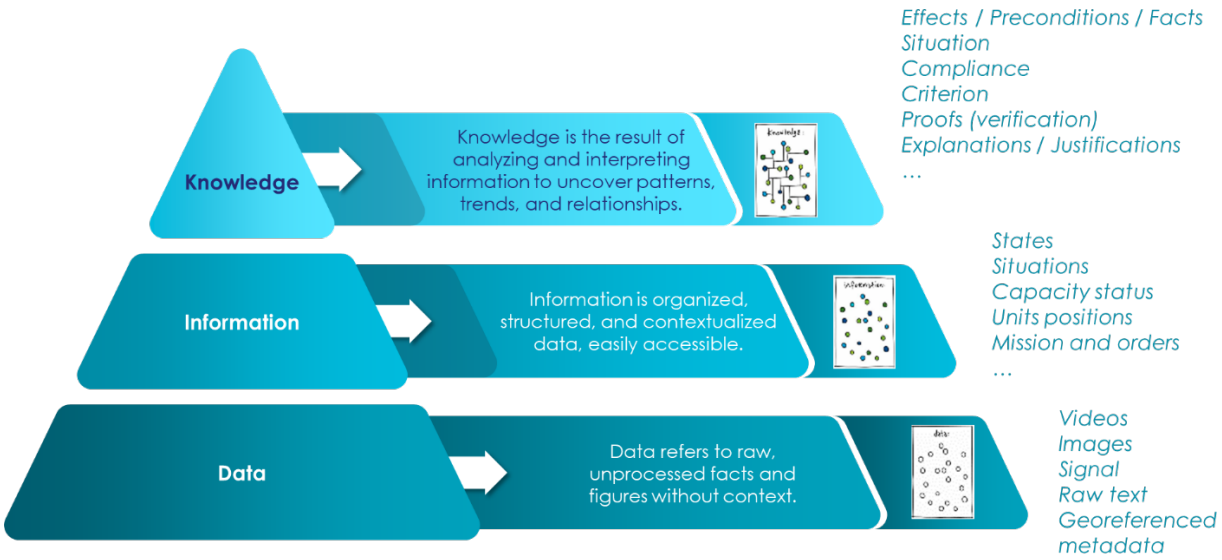

Figure 1: The data/information/knowledge pyramid.

However, information can be communicated without becoming knowledge, because its final interpretation may be performed by a human operator, as is the case in most land command systems. In such cases, information must be accompanied by its reference framework, because that framework is not implicitly shared. A body of knowledge is therefore generally specific to a profession or domain, and its meaning is shared only among domain experts. It corresponds to an ontology: a structured set of concepts that gives meaning to information.

2. What Usages for AI in Land Operations?

The contribution of AI-based solutions to land operations may be found in a wide variety of systems and capabilities. As early as 2019, the French Ministry for the Armed Forces had identified the following seven promising capabilities: (1) decision support for mission planning and execution; (2) collaborative combat; (3) cyber defense and influence; (4) logistics and operational maintenance; (5) intelligence; (6) robotics and autonomy; and (7) support, including logistic and health. These capabilities are described in the September 2019 AI Task Force report, *Artificial Intelligence in Support of Defense*.

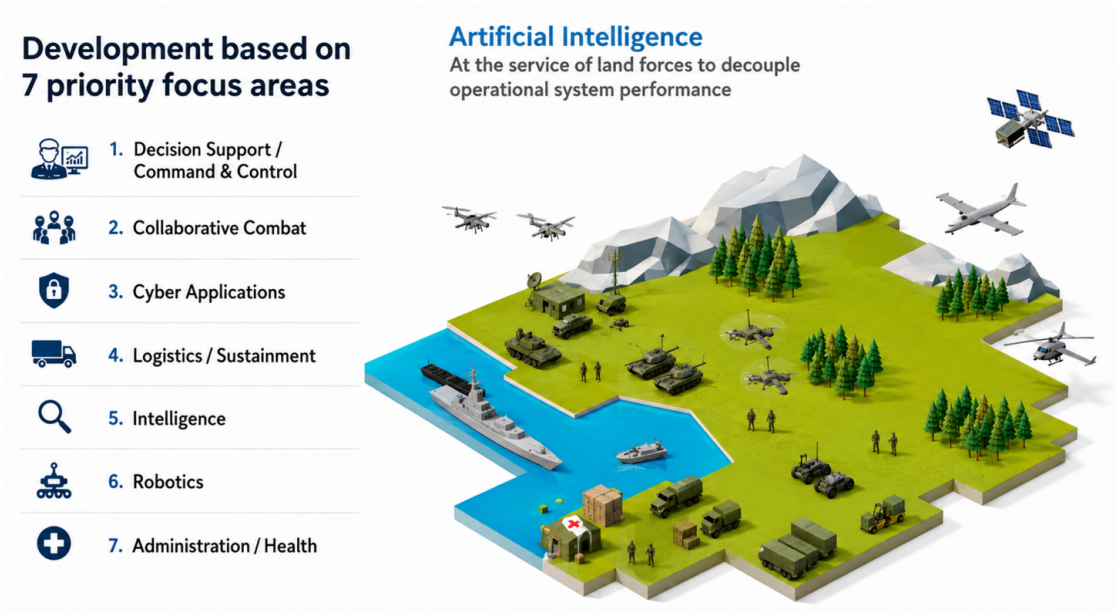

Figure 2: Operational capabilities that can benefit from AI. Image from the French Ministry for the Armed Forces AI Task Force report of September 2019

The purpose here is to illustrate some of the most striking use cases and/or those that have reached a satisfactory level of technical maturity.

Today, the military context has deeply changed. The outcome is an unprecedented density and velocity of threats, where high-intensity symmetric conflicts coexist with asymmetric and hybrid engagements. The emergence of low-cost, expendable, and attritable systems–drones, missiles, or improvised explosive devices, often assembled from civilian components–has leveled power balances and made the traditional technological superiority of Western armed forces less decisive. Following mass deployment, these tools turn battlefields into saturated spaces where engineering warfare prevails: cycles of development and countermeasure are accelerating, sometimes within a few months in Ukraine, demanding unprecedented adaptability in doctrines and industrial capabilities.

At the same time, some actors disregard ethical rules, including through automated targeting or indiscriminate attacks, while the perceived impunity of autonomous systems complicates the attribution of responsibility and blurs the red lines of international law. Faced with this loss of informational control–jamming and sensor saturation–artificial intelligence is emerging as a crucial compensatory capability, accelerating OODA loops and offsetting strategic disorientation.

Figure 3: Artificial intelligence for land operations

2.1. Radar-Based Detection and Identification of Aerial Threats

The detection and identification of aerial threats are essential to ensuring airspace security. In a world where threats evolve rapidly, it is crucial to discriminate simultaneously among many targets in complex environments. Surface radars and command-and-control centers play a key role by providing real-time analysis of the air situation, enabling rapid and accurate threat detection.

2.1.1. Relevant Capabilities

Aerial threats have evolved considerably over the years. Numerous drones and saturation attacks represent major challenges for air security. Drones in particular present unique challenges because of their small size, low radar signature, and ability to operate at low altitude. Saturation attacks, in which many objects are launched simultaneously, further complicate detection and identification. Stealth objects, helicopters masked by terrain, rockets, artillery rounds, and mortar fire are all threats that require advanced solutions if they are to be detected and neutralized effectively.

Surface radars are essential tools for detecting and identifying aerial threats. They make it possible to monitor large geographic areas and provide precise information about a target's position, speed, and trajectory. Thales surface radars, for example, are designed to detect medium- and long-range threats, from 250 km to more than 500 km. These radars are equipped with advanced software supporting next-generation aerial threats and thus provide enhanced protection against drones, stealth objects, and other modern threats.

2.1.2. What type of AI Is Used and How It Contributes

AI plays a crucial role in simplifying radar-operator decision-making. By imitating expert knowledge, statistical and connectionist AI can process massive quantities of data in real time, thereby facilitating the identification of potential threats. For example, AI can help discriminate a drone from among various objects moving at low speed, such as birds.

Deep-learning algorithms are particularly effective under unfavorable environmental conditions, where the presence of numerous unwanted detections can complicate the task.

2.1.3. Added Value

These algorithms improve drone discrimination by focusing on relevant targets. This reduces false alarms and improves decision-making in the presence of real threats. AI-based algorithms embedded in Thales radars are reliable and explainable, thereby ensuring optimal performance in critical situations. This technology also enables radar capabilities to be updated remotely in a cyber-secure manner, ensuring continuous protection against evolving threats.

Figure 4: Automatic target detection by deep learning integrated onboard a drone. Image not reproduced or translated in this draft.

2.1.4. Bringing the Solution to Its Full Potential

A collaborative network of fixed, mobile, and airborne sensors would provide comprehensive coverage without blind spots. Data fusion from different sensors –radar, optronics, and acoustics–would reduce false alarms, while analysis of micro-Doppler signatures would allow precise identification of threats such as drones through their characteristic motion patterns. In addition, a direct link to countermeasure systems, such as jammers or lasers, would enable an automated and rapid response once a threat has been confirmed.

2.2. Optronic Detection, Classification, and Threat Processing

Faced with the emergence of new threats and their continuous evolution–hypersonic missiles, combat drones, drone swarms–optronic sensors have never been more important for forces in combat, requiring an exceptional level of performance. Optronic equipment addresses comparable needs in a combat sphere closer to friendly forces. Optronic equipment such as turrets, sights, and episcopes deployed for the protection of land assets–armored vehicles and command posts, for example–also contributes to faster perception of the environment and threat detection, enabling immediate self-protection actions and real-time sharing of the tactical situation.

Figure 5: The Bertin Technologies CamSight AI is a low-SWaP camera performing real-time AI-based recognition and identification of targets–pedestrians and vehicles–in thermal imagery, with power consumption on the order of 4 W.

2.2.1. Relevant Capabilities

The development of fast drones and remotely operated munitions, with speeds on the order of 400 km/h for light vectors, increases the threat to friendly forces. This leaves at most 30 to 45 seconds between the appearance of a threat above the horizon in open terrain and a potential strike. Force protection therefore requires rapid detection, combined with an ability to discriminate what may constitute a threat, identify it, and process it under critical time constraints.

Optronic systems provide dismounted land forces with a decisive advantage on the ground by day and by night. Optronic systems complement the forward detection performed by radars with local detection close to friendly assets. They make it possible, under time constraints, to perform perimeter and sector surveillance of threats, together with classification and discrimination capabilities, by relying on long-range optronic equipment such as Safran's PASEO sights and Euroflir turrets; XTRAIM thermal weapon sights for decamouflage and threat targeting; SOPHIE portable thermal imagers for long-range detection and identification; and Thales NightRise night-vision goggles.

Figure 6: Through optronics combined with AI-based analysis techniques, the combatant will ultimately benefit from valuable assistance in perceiving the environment.

2.2.2. What AI Is Used and How It Contributes

AI-based detection, recognition, and identification (DRI) models are used to detect and classify objects of military interest in a data stream. They take in the native sensor stream, process it, enrich it with metadata, and feed it back into an exploitation environment. They can therefore serve a human user–in the case of sensors and systems directly operated by a human–or act as producers of input data for cognition and situational-knowledge modules in automated and/or autonomous systems.

Figure 7: Bertin Technologies has developed a neural network for early detection of aerial drones, enabling real-time AI detection through the analysis of background-scene anomalies while remaining insensitive to moving distractors such as clouds.

Most of these models are built through deep learning on annotated data corpora and are integrated into software solutions that enable automatic, real-time ingestion and processing of data, followed by the presentation of model predictions through various possible interfaces: the sensor-operator console, a tactical tablet controlling a drone, a command-and-control-system user interface, and so forth. These end-to-end processing chains, combining connectionist AI and expert systems, are therefore akin to neuro-symbolic AI.

2.2.3. Added Value

When integrated into PASEO sights, AI accelerates operator action and reduces cognitive load. It provides panoramic surveillance and alerting upon detection without distraction or fatigue. It then classifies observed objects, discriminating the essential–a drone or swarm–from the incidental, such as a flock of birds. Finally, it optimizes the effectiveness of the response when self-defense effectors, such as cannons and onboard weapons, are slaved to the long-range sight. In this respect, AI complements and assists the human operator by strengthening the continuity, speed, and effectiveness of action.

AI-based detection and classification tools for optronic streams have already been developed by Safran to adapt them to its ranges of airborne and land sights and initially ensure detection and classification functions, with the possible later coupling of weapons to an automatic tracker.

2.2.4. Bringing the Solution to Its Full Potential

Existing AI solutions serving this type of use case are already, for some of them, fairly mature. To go further and develop their full potential, the lines of action are clear:

- To address the scarcity of real data, develop state and industrial databases and support research on data-frugal algorithms.
- To facilitate the embedding of these functions in vectors constrained by size, weight, and available power, develop algorithms that are frugal in computing resources and define shared hardware-architecture principles.
- To facilitate deployment at limited cost across a variety of sensors and vectors, define shared principles for integration into systems and/or human-machine interfaces, including interfaces and input/output formats.

2.3. Establishing a Tactical Situation from ISR Drones

To establish a tactical situation (Tactical Situation Awareness) from ISR–Intelligence, Surveillance, Reconnaissance–sensors that may be carried by drones, and from their detection algorithms, crews face a deluge of data and information that must be fused, prioritized, and contextualized in order to provide a genuine rapid-decision capability while enabling human verification. The ability to represent the battlefield in a complete, accurate, and timely way therefore determines the relevance of military action.

In the compartmentalized, contested, and denied spaces that characterize high intensity, the ability to deploy resilient and capable tactical and/or theater intelligence assets is also a clear factor of operational superiority. It ensures the ability to observe remotely and continuously, either to alert in the event of an anomaly or threat, or to capitalize information in order to build an intelligence base on adversary patterns of life and modes of operation.[2]

2.3.1. Relevant Capabilities

Over the last 20 years, the deployment of drones dedicated to ISR or ISTAR–ISR plus Target Acquisition–missions has proven to be a way to acquire and maintain a decisive advantage in asymmetric conflicts. In conflict theaters, increased attrition of aerial vectors has not called into question the relevance of ISR assets for air-land combat support. Rather, a transition has occurred toward lighter, less costly, more versatile vectors capable of operating under electromagnetic-denial conditions.

ISR drone systems can now deploy AI capabilities to process the data streams generated by their payloads. A variety of solutions already exists, with heterogeneous technical maturity depending on whether they serve flight or mission needs. These may include resilient navigation systems–based on vision, magnetic anomalies, and other signals–that allow drones to continue their mission in the absence of reliable satellite positioning. They may include AI solutions applied to electromagnetic warfare, processing data from radio-frequency sensors and enabling the localization of adversary emitters such as radars or jammers, for example by cross-correlating direction-finding measurements. Finally, they may include AI-based DRI tools, such as Safran.AI's ODIN family of solutions for optical sensors, visible and infrared, which process sensor video streams on the ground or onboard to detect and locate military observables of interest.

The challenge of mass and the need for constant adaptation of hardware to the operational context call for versatile swarm-operation solutions in terms of both mission and platforms. AI plays a crucial role in this ability to adapt the asset to the context. This is the logic behind Thales's SwarmMaster® product, which can operate heterogeneous drone swarms in a unified way for the operator.

2.3.2. What AI Is Used and How It Contributes

For DRI and localization functions on optronic payload streams, the AI solutions are fairly similar to those described in the previous use case, involving optronic sensor-stream processing. They add another dimension: the ability to converge AI-processed streams from multiple sensors in order to build and feed a shared tactical situation as input to C2 systems at the tactical or operational levels. This includes the ability to fuse streams from several channels of the same sensor, for confirmation or decoy rejection; to converge views from multiple sensors, for confirmation or refinement of position and classification; and finally to feed the shared tactical situation with such "distilled" information using AI.

For AI solutions applied to electromagnetic-sensor streams–electromagnetic intelligence and electronic warfare–the principle is fairly similar. The objective is to detect, identify, and classify emissions of particular interest in the electromagnetic spectrum and adapt quickly, for example to activate systems such as target designation, countermeasures, or adaptive emissions. From the standpoint of artificial intelligence and software integration, the principles are similar even though the nature of the input data differs. For denied navigation solutions, particularly vision-based navigation, the capability complements optronic or electromagnetic payload-stream processing by allowing the vector to localize itself at all times, including when the integrity of the GNSS connection is compromised. The design and training of such solutions again rely on the combination of AI models trained specifically for that purpose with flexible software modules, all ultimately integrated into the drone's mission systems and flight controls.

All these elementary software building blocks are integrated to feed situational knowledge. To do so, software and AI modules are interfaced so that their combination provides a complete, augmented, and adaptive view of the tactical situation.

[2] "Patterns of life" refer to recurring behaviors and observable routines of individuals, groups, or infrastructures–movements, routines, communications–used to detect anomalies or hostile intentions.

2.3.3. Added Value

These tools have a dual purpose. For an individual vector, they make it possible both to accelerate targeting and the intelligence-to-fires loop, even in a degraded electromagnetic environment, and to contribute relevant plots for establishing the tactical situation. In the context of multiple-vector deployment, they enable collaborative situation maintenance, multi-sensor decoy rejection, the sharing of operationally relevant intelligence, and the establishment of a Common Operating Picture that facilitates rapid and appropriate command decision-making.

2.3.4. Bringing the Solution to Its Full Potential

The development of individual solutions–elementary building blocks–is progressing rapidly and has accelerated significantly across industry over the last two years, both in France and abroad. The main challenge today is to develop software means of converging these solutions so that they contribute to a complete effect for drone systems and command.

Work launched in 2025 under Project Pendragon, which includes a development dimension–led by AMIAD in connection with industry–for a C2 system serving a fully drone-enabled unit. It raises hopes for substantial progress in creating continuity between elementary AI building blocks and the system functions supported.

Semantic information fusion is becoming indispensable to build a relevant tactical situation and provide means for constructing that situation–sense-making. Since 2006, Thales has developed a hybrid AI approach combining AI-based information-aggregation techniques based on symbolic AI, such as conceptual graphs or ontologies capturing domain knowledge, with connectionist AI approaches. This makes it possible to fuse detections and tracks from several UAVs, but also to manage drones in order to ensure mission coherence, including assignment to reconnaissance or tracking tasks and drone handover management.

It will nevertheless be necessary to: (1) bring this “systems binder” to maturity within the Pendragon framework; and (2) enable the DITB to go further by adopting shared architecture principles, allowing all industrial players to develop autonomous systems combining these building blocks from an interoperability perspective. In addition, the development of low-cost, expendable drones leads to a reduction in onboard processing and therefore to a shift toward ground-based processing through a reusable C2 ground station.

2.4. Decision-Making in Classified Environments

Modern armed forces have a unique ability to capture and transmit data, the volume of which is increasing exponentially. This abundance of information can be both an asset and a challenge. It is an asset because it provides a wealth of valuable information for decision-making. However, it can also conceal the crucial data that command needs in order to make informed decisions. The massive quantity of data can make it difficult to identify the most relevant information.

Thus, data-centric C2 systems with an architecture capable of integrating AI can accelerate the cycle of the Tactical Operational Decision-Making Method, or METOD. AI can increase the precision of terrain analysis, improve the effectiveness of operation planning, and support the establishment of response plans by optimizing the allocation of assets. One example is the HexaForce C2, deployed on infrastructure of the hosting directorates based on an ARTEMIS.IA-type foundation.

2.4.1. Relevant Capabilities

Increasing the number of personnel in command centers to process this mass of data provides only limited gains. It imposes a compromise between the quantity of information to process, the quality of analysis, and the speed of decision-making. In this context, AI provides key support for the armed forces. By fusing large quantities of data from multiple sources and domains, AI amplifies human capabilities. It enables rapid crisis analysis and informed decision-making by helping identify critical information within the mass of available data.

AI aims to:

- Reduce the decision cycle from 24 hours to a few minutes by minimizing manual tasks such as note preparation, information fusion, and search.
- Process 100 times more information with the same staffing level.
- Accelerate operator training, reducing learning time by 30 percent.

2.4.2. Relevant Systems

ARTEMIS.IA–Architecture for Massive Multi-Source Information Processing and Exploitation and Artificial Intelligence–developed by ATHEA,[3] aims to provide a "sovereign and secure solution for massive data processing and artificial intelligence." This platform is intended to collect, store, cross-correlate, in real time and securely, data from multiple sources and, at the same time, provide the French armed forces with a sovereign storage infostructure.

ANTICIPE, developed by Thales, is a decision-support assistant that combines the understanding of data from multiple heterogeneous sources with collaborative recommendations involving users. It automates management of critical information–Commander's Critical Information Requirements and information processes–and Priority Intelligence Requirements. It also uses AI to inform human decisions in uncertain contexts.

2.4.3. Added Value

ANTICIPE demonstrated its effectiveness during NATO exercise Steadfast Jupiter in October 2023, where a reduced headquarters–10 operators using ANTICIPE–competed with the Brunssum headquarters of 1,000 operators. Deployed on NATO servers in Mons within a KAST classified environment, ANTICIPE was also selected by NATO ACT as a use case for its AI strategy under the decision-making pillar.

2.4.4. Bringing the Solution to Its Full Potential

Generative AI can be used to draft summaries that transmit the right information. Today, however, these techniques are not always valid–generative AI can hallucinate. It is therefore essential to ensure that automatically generated summaries are operationally correct.

2.5. Enhanced Decision-Making and Reduction of Cognitive Load

DigitalCrew® is an advanced suite of algorithms implemented by Thales and designed to reduce end-user cognitive load through automatic threat detection, classification, and tracking. As modern battlefields become saturated, sensors more complex, and threat environments more demanding, DigitalCrew® helps operators process sensor data effectively. This facilitates faster decision-making, reduces operator workload, and ultimately improves maneuverability, survivability, and lethality.

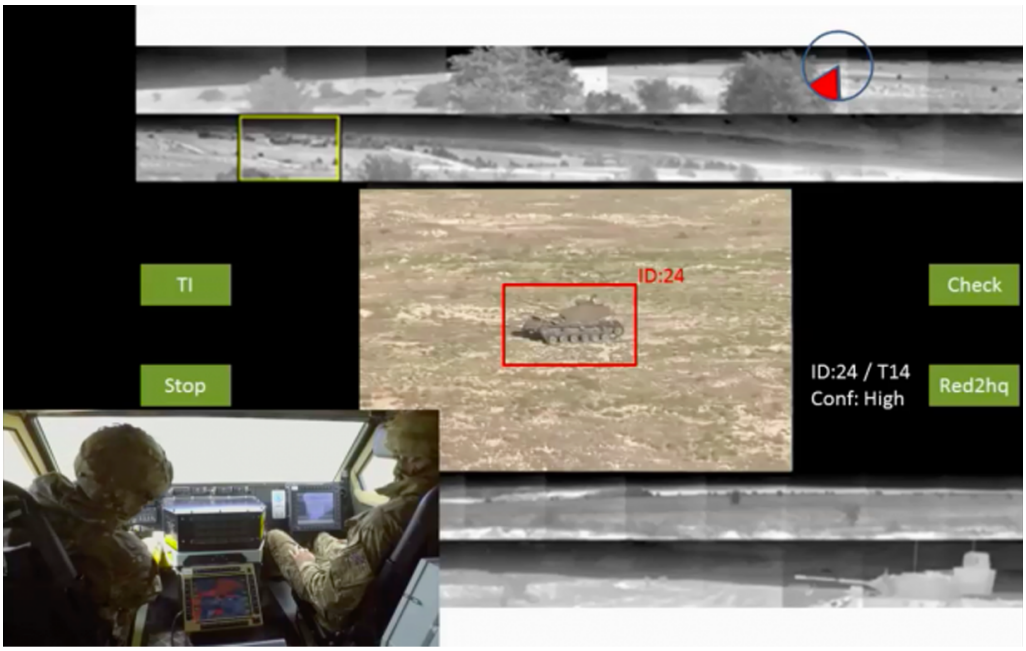

Figure 8: Once connected to one or more EO/IR sensors, DigitalCrew® reduces the decision loop by continuously analyzing the vehicle's near and far environment.

[3] ATHEA relies on an ecosystem of major industrial and digital companies, including Capgemini, Sopra Steria Group, and Airbus Defense & Space, as well as mid-sized companies, SMEs, start-ups, scale-ups, and research organizations specializing in massive data processing and AI.

2.5.1. Relevant Capabilities

User-assistance functions, such as the tracking of small and agile objects, can be deployed through simple software updates to the host platform. DigitalCrew® is versatile and independent of platform or domain. It combines traditional and machine-learning methods, including object detection, tracking, and classification; video combination; turbulence correction; and customized mission-support tools, all designed to assist operators in decision-making.

Drawing on more than a century of expertise in designing optical technologies optimized for human visual processing, Thales has now developed DigitalCrew® algorithms to work alongside human operators in interpreting and processing battlefield imagery. Unlike human operators, this digital crew member does not suffer from fatigue or distraction and maintains constant vigilance.

2.5.2. Relevant Systems

As a true example of edge AI, DigitalCrew® relies on deep expertise in hardware and artificial intelligence to enable computer vision in some of the most demanding environments, whether on land, underwater, or in the air. Its algorithms can be embedded directly on platform sensors, enabling low-latency image processing and delivering images of the highest quality.

DigitalCrew® is currently installed on unmanned demining platforms in the test phase. Some components are already deployed on the British Ajax vehicle, with integrations planned in the near future on the British Challenger 3 and on the German Joint Fire Support Team PAAG sighting system. It also supports broader counter-UAS applications.

2.5.3. Added Value

Beyond military uses, DigitalCrew® is also used in civil contexts, such as object classification during anti-poaching operations conducted from fixed-wing aircraft in Botswana.

2.5.4. Bringing the Solution to Its Full Potential

To improve the performance and capabilities of DigitalCrew®, several strategic axes can be explored, drawing on its modular architecture and advanced AI integration. First, optimizing existing algorithms is a priority. By leveraging more effective deep-learning techniques, such as convolutional neural networks or transformers, it is possible to improve the accuracy of threat detection, tracking, and classification, even in complex or saturated environments. The integration of continuous self-learning mechanisms would also enable DigitalCrew® to adapt in real time to new threats or changes in operational scenarios, thereby reducing the need for manual intervention during updates.

2.6. Automated Vehicles and Robotics

The progressive introduction of these systems makes it possible to project military power while limiting soldiers' exposure to danger and creating a mass complement. This was the purpose of the FURIOUS program, a multi-platform autonomous-system demonstrator developed by Safran Electronics & Defense and evaluated in 2022, 2023, and 2025 under realistic operational conditions. It is also the purpose of the MOBILEX challenge, led by the French Defense Innovation Agency, in which robot behavior is strongly linked to the automatic interpretation of its environment. The DROIDE program, conducted jointly by KNDS and Safran Electronics & Defense, is expected to finalize the development of a complete operational demonstrator.

Automation of land platforms is a major revolution, and numerous examples can be cited that help preserve force capital:

- Reconnaissance robots for route opening and exploration of hostile or contaminated environments without human risk.
- Demining robots for locating and neutralizing improvised charges or mines.
- Troop resupply across dangerous areas.
- Robotic weapon systems: remotely operated or semi-autonomous platforms capable of providing fire support while keeping the human in the decision loop.

The use case presented below concerns autonomous convoying developed by Arquus:

- Autonomous logistics convoys to reduce personnel risk during supply missions in dangerous areas.

Arquus has a teleoperated robotic platform offering performance advantages, particularly in terms of agility and mobility. Arquus's challenge is to equip this platform with artificial-intelligence building blocks that will ultimately confer autonomy on the system.

Figure 9: Phobos robot developed by KNDS for robotics programs.

Figure 10: FURIOUS robotic vehicle after an obstacle-crossing phase in full autonomy during operational trials at CENZUB.

2.6.1. Relevant Capabilities

One of the key functionalities to be conferred on platforms is the ability to move within the theater of operations while limiting human intervention as much as possible. This requires the robot to move over complex, unstructured, or partially degraded terrain, as opposed to paved roads. In this respect, Arquus is working in particular on environmental interpretation by its robotic platforms in order to provide them with a path-detection and path-following capability.

2.6.2. What AI Is Used and How It Contributes

To achieve the desired path-following capability, the AI must be trained on a large data bank. The AI implemented can rely on neural networks, which are particularly well suited to image analysis. In the first phase of work aimed at providing the AI with the desired path-following capability, Arquus used a database of several thousand images to train the model.

2.6.3. Added Value

Training AI on a sufficiently large data bank makes it possible to give the robot the ability to follow a detected path. Combined with the capability to interpret its environment, the robot can move autonomously across semi-structured or unstructured terrain.

2.6.4. Bringing the Solution to Its Full Potential

One line of work consists in expanding the data bank used to train the AI beyond the several-thousand-image database already used.

A second line of work consists in improving the robot's ability to interpret its environment, in particular by enabling it to infer trafficability–the quality of a terrain for being crossed by a given vehicle–which is a prerequisite for moving through a given terrain without stopping or suffering damage.

2.7. Support for Mission Planning

In modern operations –whether military interventions, the securing of sensitive areas, or humanitarian support– tactical superiority often depends on the ability of multiple units to move, act, and coordinate effectively in a complex environment. The concept of networked operations, related to the United States concept of Network Centric Warfare and the United Kingdom concept of Network Enabled Capabilities, emphasizes the quality of information sharing within the chain of command, enabling a common understanding of the situation and faster, more realistic operational planning. These concepts have now been broadened to hybrid warfare and multi-domain operations.

2.7.1. Relevant Capabilities

Translating this vision into concrete actions requires solving difficult problems: planning movements, synchronizing very different units, avoiding conflicts of use on the ground, and guaranteeing that certain actions begin only if others have been completed.

Several efforts, such as ORTAC–Operational Resource and Tactical Action Control–from Safran Electronics & Defense,[4] based on constraint programming and evaluated during Phoenix trials in 2007 and 2008, then at C4ISR-OTM in 2012 and iMUGS in 2022, with recent evolutions, propose a method for managing several tactical agents simultaneously through constraint programming and hybrid solving techniques. The environment is modeled as a graph connecting possible positions, and unit routes are then modeled as paths through a network. Each tactical unit, or agent, must not only find a feasible route but also perform actions at certain locations, respect durations, account for its capabilities, and sometimes cooperate with other units. The strength of the approach is to bring these elements–movement, actions, resources, and coordination–together in a single model, thereby producing coherent and realistic solutions.

2.7.2. What AI Is Used and How It Contributes

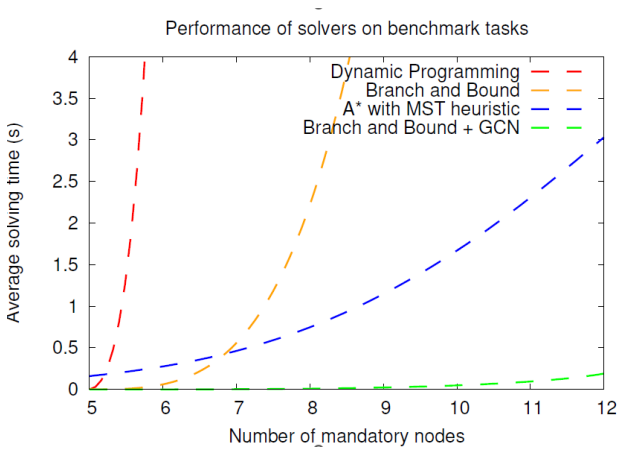

One strategy evaluates search "forward" across different routes while satisfying coordination constraints as closely as possible. The neuro-symbolic technique used makes it possible to find realistic plans more quickly and sometimes even prove that they are optimal. Recent work shows the value of graph-learning (GNN) techniques for characterizing the structure of the problem. This improves solving, whether using neural networks or symbolic learning, such as the dynamic learning of logical clauses. The figure in the source document shows the performance obtained in green by a hybrid neuro-symbolic approach compared with classical methods on a realistic problem, as presented at IROS 2019. Other approaches have also shown the value of managing conflicting constraints derived from Conflict-Based Search. These approaches combine scalable performance with better explainability of the solutions obtained.

2.7.3. Added Value

ORTAC has been evaluated several times in a Battle Lab across representative scenarios: reconnaissance in urban areas, reinforcement of allied positions, inspection of sensitive sites, and securing humanitarian zones. Several efforts also adapt the approach to autonomous agents.[5] The results show that the method can generate complex plans for several agents, possibly robotic, while respecting operational constraints. They also show that

[4] See *Solving Planning and Scheduling Problems in Network-Based Operations*.

[5] The European iMUGS project, followed by PARHERO conducted with ONERA and completed in 2025.

certain algorithms–notably those guided by heuristics and evaluation functions–provide significant gains in computation time.

2.8. Predictive Maintenance and Operational-Readiness Support

Predictive maintenance is much more than a simple technological evolution; it embodies a profound transformation in the management of critical systems. Traditionally, maintenance strategies relied either on corrective interventions triggered after a failure or on preventive approaches planned according to fixed schedules or predefined wear thresholds. While effective in some contexts, these methods have major limitations: they do not always account for real operating conditions or environmental and operational variability. This can lead either to premature maintenance, generating unnecessary costs, or to late maintenance, causing costly breakdowns and increased risk.

2.8.1. Relevant Capabilities

In this context, predictive maintenance emerges as a strategic capability that anticipates failures before they occur. Its fundamental principle is based on early detection of even minute signs of degradation and on the precise estimation of the time remaining before failure, commonly called Remaining Useful Life. This approach does not merely predict when a failure might occur; it also seeks to understand why, by identifying the root causes of observed anomalies. It therefore makes it possible to shift from a logic of repair to a logic of active prevention, optimizing both system availability and maintenance costs.

This capability is particularly crucial in sectors where reliability is non-negotiable, such as military operations. An unexpected failure in such domains can have consequences much more serious than a mere production stoppage: it can affect personal safety or mission continuity. Predictive maintenance therefore provides a response suited to these challenges by aligning interventions with the actual condition of equipment rather than with theoretical assumptions.

2.8.2. What AI Is Used and How It Contributes

To implement effective predictive maintenance, a hybrid AI approach can be adopted, combining the strengths of symbolic and connectionist AI. Symbolic AI, based on formal models such as FMECA or fault trees, uses expert knowledge to analyze well-documented critical systems. It models the nominal behavior of equipment and detects deviations indicative of faults, drawing on decades of operational experience. This method is particularly suited to environments in which degradation mechanisms are mastered, providing the explainability and rigor essential to regulated sectors such as aerospace and defense.

Conversely, connectionist AI excels at analyzing massive and complex data from sensors or logs. Neural networks identify degradation patterns in time series, while techniques such as hidden Markov models estimate future system states. This approach is ideal for poorly documented systems where raw data is the main source of information.

The hybridization of these two AI approaches creates optimal synergy: physical models validate algorithmic predictions, while data enriches expert rules. The result is a set of adaptable solutions ranging from standardized equipment to the most complex dynamic systems. Finally, generative AI complements the approach by exploiting maintenance reports to guide technicians and automate report writing, thereby strengthening overall efficiency.

2.8.3. Added Value

The adoption of AI for predictive maintenance is not limited to technical optimization; it generates global added value across economic, operational, and strategic dimensions.

Economically, the gains are immediate and measurable. By avoiding unexpected failures, predictive maintenance significantly reduces costs associated with emergency repairs while extending equipment life. By optimizing the planning of interventions, it also limits overstocking of spare parts and unnecessary interventions, thereby freeing resources for higher-value activities.

Operationally, AI-based predictive maintenance considerably improves system availability. By anticipating failures, it makes possible to schedule maintenance stoppages during low-activity periods, thereby minimizing the impact on operations.

2.8.4. Bringing the Solution to Its Full Potential

Predictive maintenance offers clear advantages, but large-scale deployment is difficult. The first challenge is the availability, quality, and integration of data. AI algorithms depend on them. Yet data is often heterogeneous, noisy, or stored in disconnected systems. Moreover, current systems were not designed either to store or to process such data. It is therefore necessary to evolve these systems and establish a data logistics process to

capture data, develop these approaches, integrate onboard alert systems, and maintain continuous learning so that systems remain faithful to operational experience.

The second challenge is solution scalability. Pilot projects are often very useful, but generalizing them is complex. To achieve this, a modular approach is needed, with AI models that are reusable and adaptable to different contexts. For example, an algorithm trained on tank-engine data could be reused for light vehicles, provided that it is enriched with specific data. This transfer-learning strategy reduces deployment costs and timelines.

2.9. Military Logistics

Military logistics encompasses all actions aimed at supporting armed-forces operations. It is defined as the science of planning and executing the movement and maintenance of forces. It covers a wide range of domains: acquisition, maintenance, and repair of materiel and equipment; transportation of personnel, materiel, and equipment; construction and maintenance of facilities; supply of fuel, food, and ammunition; acquisition or provision of services; and medical and health support. Logistics are essential for maintaining combat capability and has therefore become a target for enemies, especially as deep-strike capabilities are increasing.

Modern operations also require logistics able to adapt rapidly to changing situations. Flexibility makes it possible to modify objectives or seize opportunities, while resilience ensures continuity of operations even in the event of partial failure. Finally, the need to maintain robust and flexible logistics is confronted with limited budgets, requiring planners to optimize resources.

AI is thus transforming military logistics by improving the planning, responsiveness, and efficiency of supply chains–factors that are decisive for operational success–while reducing risk to soldiers. Its applications cover several key domains, from route optimization to resupply in hostile areas.

2.9.1. Relevant Capabilities

AI strengthens three critical logistics pillars for land forces.

Resource planning and optimization

- Stock and flow management: anticipate needs for ammunition, fuel, food, and medical supplies in real time.
- Route optimization: select the safest and most efficient convoy routes while integrating dynamic constraints, such as enemy threats, weather conditions, and the state of infrastructure.
- Multinational coordination: synchronize resources among allies, for example by mutualizing surpluses among NATO contingents.

Autonomous and robotic execution

- Automated resupply: use autonomous vehicles and drones to deliver supplies in hostile areas, reducing soldier exposure.
- Robotic medical evacuation: transport the wounded while minimizing risk to medical personnel.
- Predictive maintenance: continuously monitor the condition of equipment–vehicles and weapons–to avoid critical failures.

Resilience and real-time adaptation

- Logistics-crisis management: react immediately to unforeseen events, such as stockouts, attacks on depots, or communications jamming.
- Degraded-mode operation: maintain operational capabilities even in the event of partial loss of AI systems, for example by returning to manual procedures.
- Proactive cybersecurity: detect and neutralize cyber threats targeting logistics chains.

2.9.2. What AI Is Used and How It Contributes

For logistics planning, connectionist AI–machine learning–and optimization algorithms analyze massive volumes of historical and real-time data. They predict future needs by cross-correlating information such as equipment wear, weather conditions, and enemy movements, then suggest optimal routes or reallocations of resources. Such systems can, for example, identify that one contingent has a surplus of medicines while another is critically short and propose automatic redistribution.

In autonomous execution, onboard AI in vehicles or drones takes over. Equipped with advanced sensors, such as cameras and LiDAR, and navigation algorithms, these systems can move through hostile terrain, avoid obstacles or threats, and deliver their payload with extreme precision, even under enemy fire.

AI also contributes to predictive maintenance: by continuously monitoring the condition of vehicles and equipment, it detects early signs of failure and triggers alerts for intervention before materiel breaks down, thereby reducing unplanned immobilization.

To guarantee resilience and cybersecurity, neural networks and anomaly-detection systems can be deployed. These tools continuously scrutinize logistics data flows to identify suspicious activity, such as attempted hacking or falsification of orders. They also simulate crisis scenarios–destruction of a depot, communications jamming– in order to assess the system's ability to respond and adjust protocols accordingly.

Finally, multicriteria decision systems can support resilience and help choose the right degraded-mode configurations, switching to fallback protocols when primary systems are compromised.

2.9.3. Added Value

The benefits of AI in military logistics are multiple. From a human standpoint, it considerably reduces soldiers' exposure to danger, particularly by limiting their participation in supply convoys, which are frequently targeted by the enemy. Operationally, it improves overall effectiveness by reducing delivery times, optimizing routes to save fuel, and avoiding stock-management errors such as overstocking or shortages. These gains translate into better availability of front-line units, which can focus on their primary mission rather than on logistics problems.

AI also strengthens operational resilience by allowing forces to maintain capabilities despite major disruptions, whether cyberattacks, physical attacks on depots, or sudden deterioration in weather conditions. Finally, AI offers a strategic advantage by facilitating coordination among allies, for example within NATO or international coalitions. By harmonizing logistics data and proposing resource mutualization, it creates collective superiority over less well-equipped adversaries.

2.9.4. Bringing the Solution to Its Full Potential

International collaboration and standardization play a key role. Armed forces must work together to establish common protocols for exchanging logistics data, avoiding incompatibilities between systems. Partnerships with the private sector, particularly companies specializing in Logistics 4.0 or cybersecurity, can accelerate innovation and the adaptation of civilian technologies to military needs.

2.10. Training and Simulation

AI is also transforming force preparation:

- Advanced virtual environments: realistic simulation of theaters of operations enabling immersive training.
- AI adversaries: creation of virtual opponents able to adapt their tactics for more realistic exercises.
- Performance analysis: objective assessment of tactical decisions and soldier reactions to identify areas for improvement.
- These tools enable training that is more intensive, more varied, and less costly than traditional field exercises.

2.10.1. Relevant Capabilities

On an increasingly complex battlefield, with actions and effects conducted simultaneously across different domains, mission preparation, assessment of operational conditions, and training are more difficult, intensive, and resource-consuming. In addition, the diversity and dynamics of the systems used no longer allow monolithic and static simulation environments to be developed, whether to simulate friendly or enemy forces. The Battleverse project, funded under the European Defence Fund, responds to the military requirement to prepare missions, examine decision-making processes, and simulate operational actions while covering as many possible scenarios as possible and considering the use of varied systems. The project involves a broad range of industrial actors, including Thales, Sopra, Safran, and Naval Group. This capability must be usable both over the long term for fictional scenarios and over the short term with realistic operational situations.

2.10.2. What AI Is Used and How It Contributes

The approach considered consists in synthesizing, on demand, digital twins of the battlefield and of the systems that compose it. The project considers the generation of scenario data or multiphysics behaviors. In this framework, the use of neuro-symbolic generative AI enabling the expression and control of physical models is envisaged. This helps compensate for the lack of data, information, or models regarding objects on the battlefield.

In addition, the development of AI techniques addresses several fundamental simulation-application problems:

- **Useful time:** accelerating the OODA cycle during mission planning and execution. AI techniques must be sufficiently interactive and return symbolic information that can be used immediately, without saturating cognitive load.
- **Explainability:** accounting for the human in the loop or on the loop, mainly for situation maintenance but also for execution control. The coherence of the information provided must be ensured.
- **Online adaptation of AI:** for example using reinforcement-learning techniques to optimize the execution of actions according to the operational context while reducing margins of error and considering strict constraints such as rules of engagement.

The whole is integrated into a flexible, scalable, and fast-executing military simulation framework for wargaming, storytelling, or immersive situation training for multi-domain operations. Simulation engines must be able to adapt to the growing complexity of modern warfare.

2.10.3. Added Value

Force preparation, operational tempo, uncertainty, and bias in the apprehension of the battlefield are all concerns addressed by this research project. The value for the chain of command is therefore undeniable, and this type of on-demand digital twin provides several advantages:

- Grasp the variable-geometry problem of the battlefield and the emergence of improvised systems, particularly those endowed with autonomy.
- Open the operational field of possibilities during design, whether for deliberate or time-sensitive planning, while enabling the evaluation of different engagement situations across a very large number of scenarios.
- Improve systems engineering by enabling a better understanding of system uses and their appropriation by forces, particularly for systems with AI or autonomous capabilities. Such digital-twin environments can act as a bridge between operational personnel and engineers.
- Benefit from the power of generative design to accelerate the engineering cycle and provide faster and smoother design-experimentation loops.

2.10.4. Bringing the Solution to Its Full Potential

The development of libraries of models, scenarios, and use cases is the main capitalization axis for such a digital-twin environment. The ability to develop new hybrid AI systems and specialize models on concrete and realistic scenarios also helps improve the robustness of digital twins. Finally, capitalizing on the data, information, and knowledge generated by the generative environment also helps concretely address a general lack of data regarding the variability of future operations.

2.11. Rapid Detection of Loitering Munitions for Soft Kill

2.11.1. Relevant Capabilities

Detecting loitering munitions, even at night, requires rapid reaction to flying objects that first appear as very small objects at long range and may arrive from any direction, from the sky down to near-ground level. A rapid automatic alert system based on infrared video can trigger a soft-kill self-protection system on an armored vehicle

in time, such as systems developed by Lacroix Défense. AI can, for example, frame the infrared image region of interest so that the user can rapidly confirm the object's identification and trigger the soft-kill system.

Figure 11: Principle of drone detection for soft-kill protection.

2.11.2. What AI Is Used and How It Contributes

Neural networks are particularly well suited to image analysis, not only in natural light–the most widespread case–but also in the infrared, as illustrated by Bertin Technologies' CamSight AI.

Integrating AI directly into the camera makes installation more practical and gives the alert system greater robustness, because the attack surface of the detection system is limited by the absence of an external computer, which could itself be damaged.

2.11.3. Added Value

Non-AI methods are typically based on detecting changes in images, such as a point moving on the horizon. This makes them more sensitive to situations such as the motion of clouds or tree leaves in the wind. Such methods then require more complex algorithms that are harder to develop in order to manage these cases.

2.11.4. Bringing the Solution to Its Full Potential

As always with machine learning, it is necessary to have sufficient data representative of the real conditions in which the loitering-munition detection system will be used: variable weather, drones with varied behaviors, varied driving conditions, and so on. This requires an image-acquisition effort typically spread over time and followed by serious annotation work.

Ensuring sufficient image-analysis speed is also important. The difficulties arise from the fact that embedded processing is typically constrained in resources–SWaP: Size, Weight, and Power.

3. What Methods?

3.1. One or Several Disciplines?

Historically, the design of AI algorithms emerged in the 1950s through two currents. Knowledge-based AI, now often called GOFAI–Good Old-Fashioned AI–or symbolic AI, is based almost exclusively on symbolic reasoning and different forms of logic.

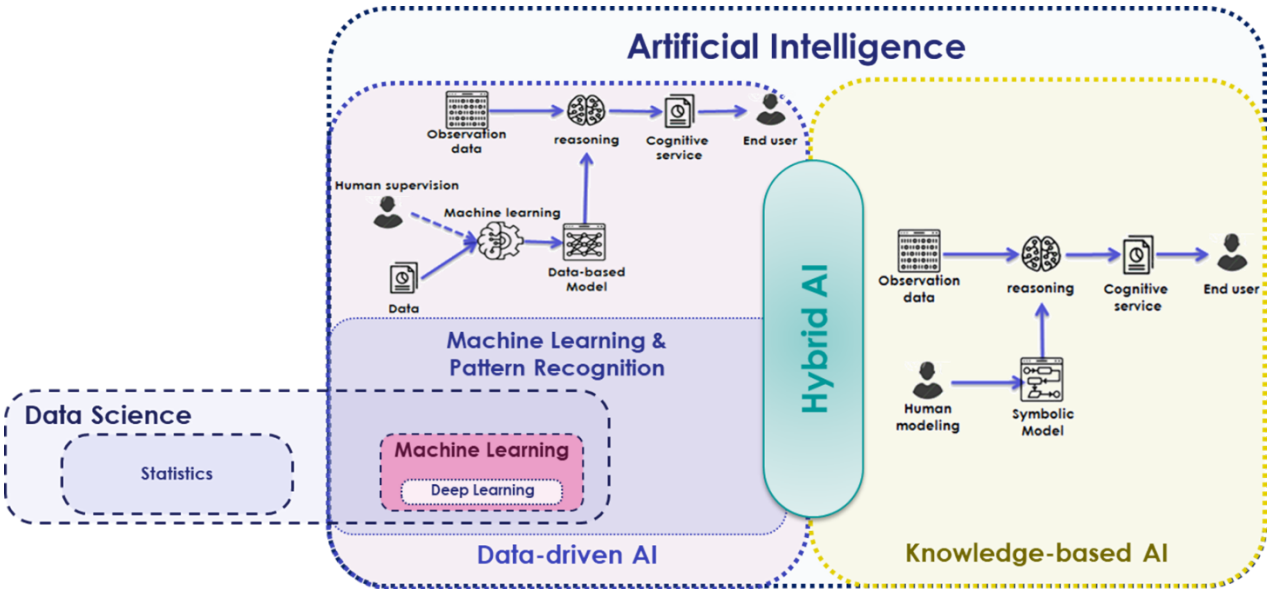

Figure 12: Main AI paradigms.

It differs from data-driven AI, also called statistical and connectionist AI, which has been in the spotlight in recent years because of massive data collection and the arrival of learning techniques such as deep learning and generative AI.

Symbolic or Knowledge-Based AI
An AI approach based on the explicit manipulation of logical rules and knowledge formalized by human experts. Reasoning is transparent and traceable. Example: during the 1991 Gulf War, the expert system DART–Dynamic Analysis and Replanning Tool–planned and replanned logistics transport for soldiers and materiel in real time. According to DARPA, within four years it saved an amount equivalent to the total AI budget invested since the 1960s.
Statistical and Connectionist AI, also known as Data-Driven AI
An approach in which the machine learns from large quantities of data by identifying correlations and statistical regularities, without being explicitly programmed with rules. Machine learning, a major statistical-AI technique, entered companies around 2010-2015. Example: an algorithm trained on thousands of past alerts to predict the probability of a threat according to context.
Generative AI
A subfamily of statistical AI able to produce new content–text, images, sounds, video, or code–by relying on models trained on very large corpora. Example: an assistant automatically drafting an operational report from voice notes or generating a realistic training scenario.

Whereas symbolic AI uses knowledge transmitted to the machine as models in order to solve problems, data-driven AI starts from solution examples that it tries to extrapolate through statistical and probabilistic methods in order to build an artificial model. Their areas of employment, uses, and application objectives therefore differ significantly.

Today, AI encompasses many disciplines, including machine learning, knowledge management, logical reasoning, problem solving, multi-agent systems, natural-language processing, intelligent robotics, and probabilistic-inference techniques. Figure 13 in the source document presents most AI technologies according to symbolic or data-driven approaches.

Machine Learning
An AI technique in which a program improves its performance on a task through examples rather than by being manually programmed. Such learning is typically performed at discrete moments and produces a system whose performance is fixed until the next retraining. Example: a fingerprint or face-recognition system associating identities with names and improved from new confirmed identifications.

Deep Learning
A machine-learning technique based on artificial neural networks with many successive layers of computation. This technique is particularly effective for processing complex data such as images, sound, or text. Example: automatic target detection in drone video streams or transcription of radio communications.

Probabilistic techniques, based on Markov processes or Bayesian inference, are often overlooked but continue to develop under the concept of the "belief state." Over time, the symbolic component has been structured into several related domains, such as logical inference, problem solving, and explainability.

Probabilistic Inference
A method enabling a system to reason under uncertainty: rather than providing a binary true-or-false answer, it estimates the probability of different hypotheses from partial or ambiguous data. Example: evaluating the probability that an observed convoy belongs to a given enemy force by combining imperfect human intelligence, imagery, and electronic signals.

The field of multi-agent systems, in which agentic AI is currently developing, provides the environment for constructing AI architectures. Agentic AI refers to the ability of a system to act autonomously and intelligently, adapting to its environment and making decisions in real time. Agents therefore have an impact on all the technologies developed and have a lasting influence on the state of the art in distributed computing based on network communications. Although so-called symbolic techniques are often opposed to connectionist and probabilistic artificial intelligence, many hybrid approaches are emerging. In particular, agentic AI is the most accessible form of hybridization. This systematic approach is now extending agent systems to language models and pushing the combination of AI techniques to its extreme.

Multi-Agent Systems
An architecture in which several autonomous entities–the agents–interact to accomplish objectives. Interactions may be cooperative, when agents coordinate toward a common goal, or competitive, when agents pursue antagonistic objectives, which can be used to strengthen their mutual capabilities. Examples: a fleet of autonomous drones that scout an area without duplication; or two agents simulating opposing forces in a digital wargame to test doctrines, tactics and procedures.
Agentic AI
AI systems able to pursue an objective by deciding for themselves, at each step, what actions to undertake and which tools to mobilize. Unlike an automated pipeline with fixed stages, an agentic system adapts its behavior according to intermediate results and can manage unforeseen situations. In practice, these systems currently often rely on large language models to provide reasoning and dynamic decision-making. Example: an agent asked to "summarize the week's security incidents" will autonomously locate relevant documents, extract key information, detect inconsistencies, and relaunch an analysis if necessary, without any of these steps having been explicitly programmed.

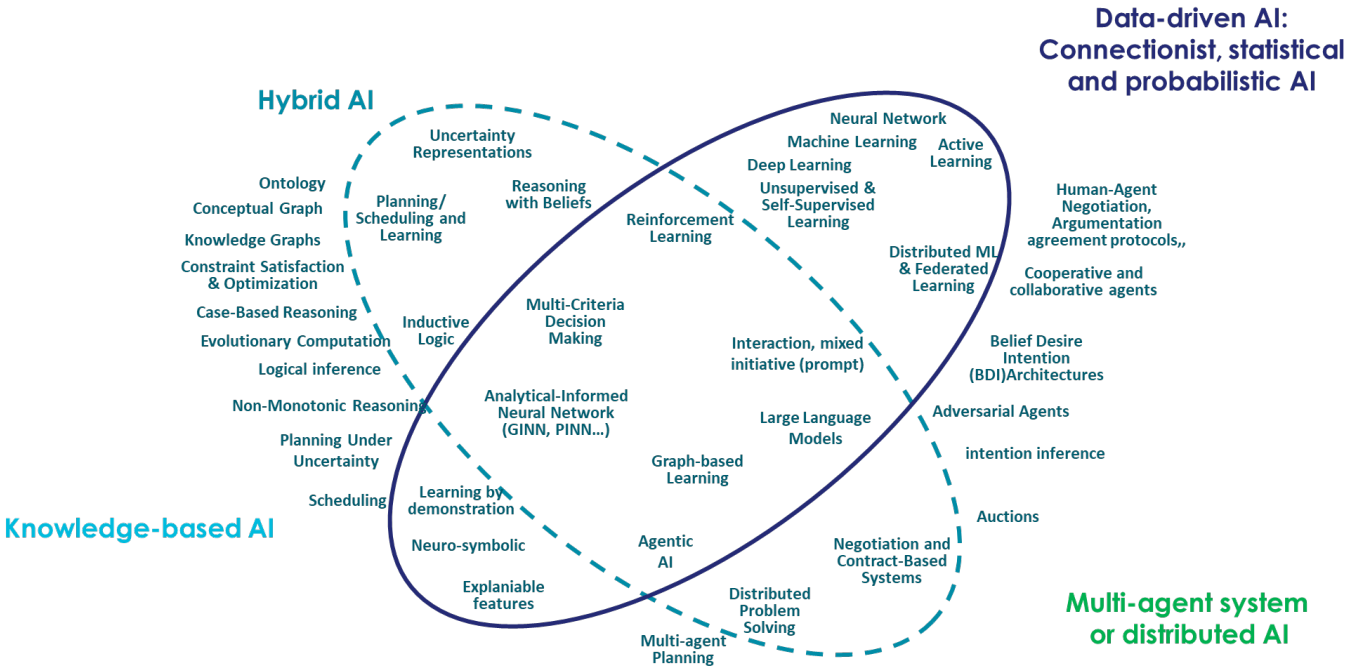

Figure 13: The zoology of AI technologies.

David Sadek, Vice President for AI and Algorithms at Thales, explains that "connectionist AI is the AI of the senses, and symbolic AI is the AI of meaning." For this reason, to cover all cognitive capabilities, the future lies in the hybridization of the two approaches.

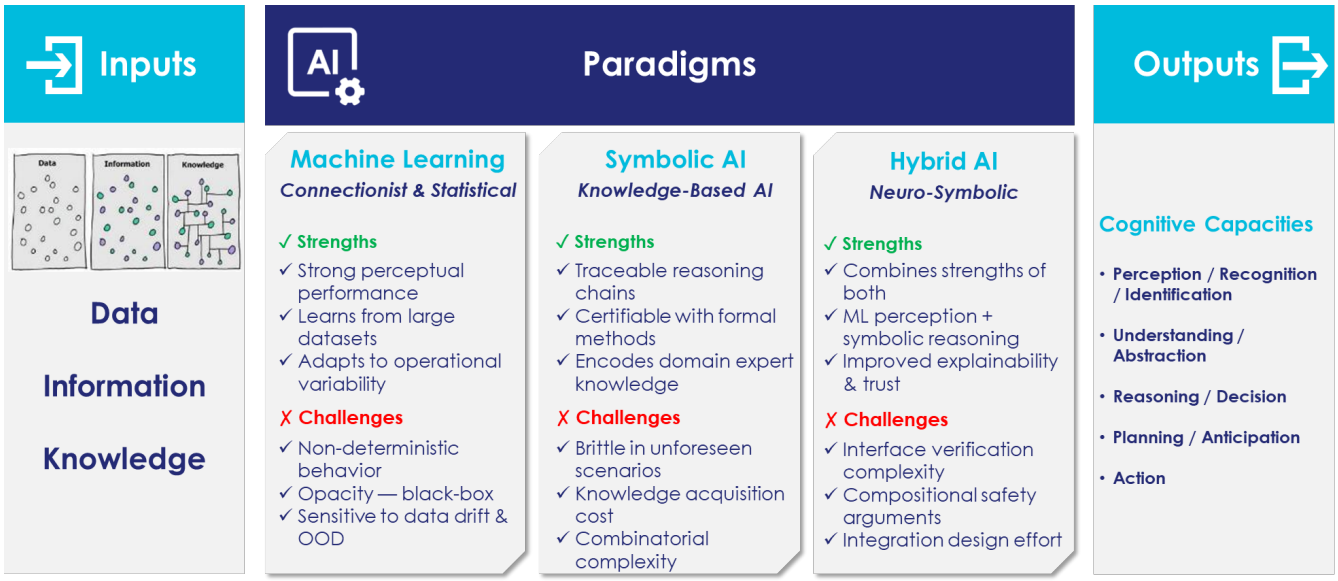

Figure 14: Strengths and weaknesses of the different AI paradigms.

3.1.1. Mapping AI for Land Operations

In an uncertain geopolitical context and faced with increasingly complex and asymmetric threats, the introduction of AI into land operations can considerably accelerate the pace of operations, sometimes redefining tactical, operational, and strategic paradigms.

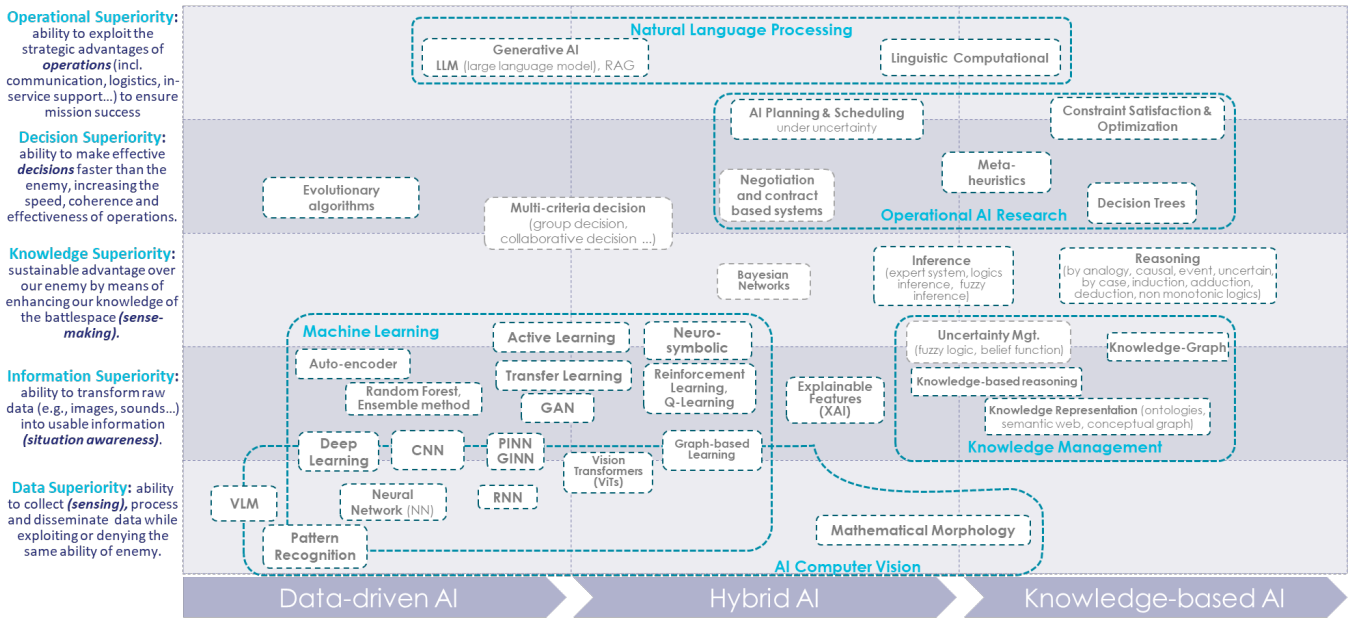

Figure 15: AI paradigms considered against their operational areas of use.

The ability to process information rapidly, make informed decisions, and execute them without delay provides command with a clear advantage, as emphasized by the use cases described in Section 2. Reducing the decision cycle makes it possible to gain a decisive advantage over a technologically less advanced adversary, who cannot react to rapid changes in the tactical situation.

This concept, initiated in the field of networked operations,[6] has now been generalized and systematized across all operational domains. It is also the foundation of modern hybrid warfare and multi-domain operations, enabling planned coordination of assets in order to deliver synchronized effects across several domains. This precursor concept anticipated many principles beneficial to the exploitation of AI: separation of planning and execution, or mission control, in a continuous, flexible, and adaptable process; collective planning and execution; reduction of the fog of war for a more transparent battlefield; and so on.

This approach enables the construction of a framework for reading AI-enabled capabilities.

3.1.2. Data Superiority

Mastery of data is a decisive asset. Through massive collection by advanced sensors–satellites, drones, radars–colossal volumes of information are generated in real time. Rapid processing transforms this raw data into actionable intelligence within seconds. This supports informed decision-making, optimizing operation planning, strike precision, and logistics effectiveness. The integration of AI into land operations therefore opens major perspectives, but its effectiveness depends closely on mastery of data and of its life cycle.

Unlike traditional software tools, new generations of models rely on statistical calculations requiring representative, voluminous, and diverse data to ensure reliable predictions under real conditions. Several challenges remain, however, including the availability of training and test data, their adequacy with operational contexts–sensor flows, ontologies, and domain standards–and their volume and variety to cover all use cases without bias.

Data sovereignty is a critical issue, especially for sensitive data such as satellite or military images, and annotation further reinforces their strategic value. Protecting such data against leaks or undesirable exploitation requires strict sharing frameworks, particularly with industrial partners, while clarifying intellectual-property rights and export restrictions. Innovative solutions, such as crowdsourcing–as in the Ukrainian example–or the use of simulated data can complement real datasets, but they raise questions of confidentiality and licensing.

Model retraining is also a major challenge, especially given the rapid evolution of operational contexts, as illustrated by the war in Ukraine. It requires robust methodologies to incorporate new field data while maintaining performance, as well as annotated test data to validate models quantitatively without sharing those resources with third parties. Finally, the physical and cyber security of data must be ensured at each stage, from acquisition to storage and processing, in order to preserve operational advantage and confidentiality.

[6] A concept attributable to David S. Alberts, *Network Centric Warfare: Developing and Leveraging Information Superiority*.

Meeting these challenges requires close collaboration among data scientists, operational personnel, and industry in order to define procedures adapted to data acquisition, annotation, sharing, and updating, while anticipating risks linked to obsolescence or declassification. The ultimate objective is to reconcile model performance, data sovereignty, and operational agility in order to make AI a decisive lever on the battlefield.

Cybersecurity also plays a key role by protecting these flows against cyber threats, thereby preserving the confidentiality of strategies and mission continuity. Finally, predictive analysis provides a major strategic advantage: anticipating adversary movements shifts the balance of power in favor of the actor that dominates information. Ultimately, military superiority rests first on the ability to collect, analyze, and secure data, thereby guaranteeing responsiveness, precision, and tactical ascendancy on the ground.

3.1.3. Information Superiority

Information superiority rests on the ability to collect, process, and analyze immense volumes of heterogeneous data in real time, requiring AI technologies that are particularly effective in several domains.

Connectionist and statistical AI–such as deep learning and neural networks–is particularly suited to the analysis of unstructured data, including satellite images, radar signals, and intercepted communications. These data-driven AI technologies excel at detecting and recognizing complex patterns and anomalies in large datasets and information sets, making it possible to identify weak signals that escape human analysis or traditional statistical methods.

Generative AI is relevant for multimodal data fusion because it can integrate information from diverse sources–images, text, acoustic and electromagnetic signals–into a coherent and exploitable representation. Generative models can also fill information gaps by producing plausible hypotheses about intelligence blind spots.

Multi-agent systems enable decentralized collection of data and information produced by intelligent and distributed sensor networks. Multi-agent architectures are particularly robust in the face of disruption and can maintain situation awareness even in the event of partial degradation of the information network.

Integrating these technologies into hybrid architectures enables domination of the information spectrum at all levels, from tactical collection to strategic analysis of adversary intentions. In particular, AI-assisted surveillance systems radically transform terrain reconnaissance and threat detection:

- Through machine-learning techniques, anomaly detection helps identify suspicious behavior or subtle environmental changes that may indicate an imminent threat. Data-driven AI can therefore automatically identify enemy vehicles, installations, and troop movements from aerial or satellite imagery.
- Hybrid AI based on combinations of physical or geometric models with neural techniques–Physics-Informed Neural Networks and Geometric-Informed Neural Networks–using historical data provides detection, recognition, and prediction capabilities that are particularly robust, especially against sensor-induced noise.
- Multi-source information fusion–radar, infrared imagery, acoustic or electromagnetic signals, and open sources–based on knowledge-graph technologies such as semantic networks, conceptual graphs, or ontologies, or on generative AI, helps build an enriched operational picture for a better understanding of the tactical, operational, and strategic situation. Such multi-source information fusion optimizes military effectiveness by providing predictive analysis and a global understanding of the operational environment.

All these technologies contribute to better situation awareness and thereby reduce the risk of ambushes or surprise attacks.

3.1.4. Knowledge Superiority

Mastery of knowledge, combining raw data and domain expertise, is an indispensable strategic lever for armed forces. Unlike the simple accumulation of data and information, knowledge superiority must integrate military doctrines, lessons learned, and operational know-how in order to transform intelligence-derived information into relevant and adapted decisions. A deep knowledge of tactical procedures, adversary capabilities, and operational environments makes it possible to refine plans, anticipate enemy reactions, and adjust actions in real time.

This superiority reduces uncertainty and risk by relying not only on updated data but also on a fine understanding of rules of engagement, decision patterns, and lessons drawn from past conflicts. It also optimizes the use of resources by aligning available means with doctrinal imperatives and field realities. Finally, it provides an asymmetric advantage: by combining technical intelligence–SIGINT and HUMINT–with tactical know-how, such as NATO doctrines and national procedures, a force can disorient the adversary, exploit weaknesses, and maintain the initiative.

Manipulating data or information can maintain the cohesion of a chain of command or manage the correctness and integrity of a platform state, whether aircraft, ship, or robot. However, even if neural-inference techniques can be developed, not all information necessarily lends itself to knowledge manipulation. Reasoning tasks such as deduction, induction, abduction, predicate satisfaction, or constraint satisfaction must rest on well-defined logical semantics, for example doctrines. The problem is obvious in decision support for a chain of command, where temporal logics must be guaranteed–far beyond usual informational representations. To benefit from AI

assistance that is more explainable, it is necessary to follow logical and temporal sequences, such as DigitalCrew® causalities. A logical explanation provided in useful time would obviously enable the chain of command to assume responsibility for a difficult decision. Another example is that the behavior of an autonomous robot or drone must be explainable to the operator in order to avoid operational questioning and hesitation.

In all cases, having an automated reasoning framework will enable assumed decision-making and, above all, anticipation relative to the adversary's decision cycle.

In short, knowledge superiority–well beyond simple information superiority–combines human expertise and structured data to provide superior decision capability, essential to the success of contemporary military operations.

3.1.5. Decision Superiority

One of AI's major contributions is the improvement of decision-making processes. AI-based decision-support systems can process considerable volumes of data in real time to present commanders with a clear view of the battlefield and contribute to the following capabilities:

- Predictive analysis to analyze enemy movements, understand intent, anticipate future actions, and suggest optimal countermeasures.
- Operational planning to rapidly evaluate multiple tactical scenarios and select the most effective plans.
- Resource management to optimize the allocation of troops, vehicles, and ammunition according to operational priorities and logistics constraints.

These capabilities accelerate the OODA loop, giving a decisive advantage to forces that master these technologies.

Decision superiority nevertheless requires capabilities for complex reasoning, anticipation, and adaptation to novel situations, calling on different but complementary AI technologies.

Symbolic AI technologies are particularly relevant in domains where formal logic and legal or doctrinal constraints must be rigorously respected. Such systems enable explicit reasoning that complies with doctrines, rules of engagement, the law of armed conflict, and ethical considerations framing the use of force.

Hybrid AI is the most promising approach for decision superiority, combining the flexibility of data-driven AI, such as learning, with the rigor and explainability of symbolic or knowledge-based approaches. It offers an optimal balance between adaptation to a dynamic and uncertain context and compliance with doctrinal and regulatory frameworks.

Finally, simulation and digital twins are essential for anticipating the consequences of decisions, making it possible to virtually explore multiple tactical options and their probable effects.

Unlike symbolic AI approaches such as constraint programming, which are effective for designing optimized plans under resource constraints, generative AI can design entirely new maneuver schemes adapted to situations evolving in a dynamic and uncertain future. It can therefore propose innovative operational plans by exploring unconventional solution spaces.

The complementarity of these approaches makes it possible to create decision-support systems that amplify the cognitive capabilities of commanders while keeping the human at the center of the final decision-making process.

3.1.6. Operational Superiority

Operational superiority concerns the optimal execution of actions on the ground, requiring AI technologies capable of operating in real time in complex and dynamic physical environments.

Reinforcement learning is particularly suited to the automated control of robotic platforms and continuous optimization of tactics based on field feedback. It enables systems to adapt rapidly to changing operational conditions and improve performance over successive missions.

Collaborative multi-agent systems are becoming indispensable for operations involving swarms of drones or ground robots, as well as coordination among human and robotic units. These architectures allow sophisticated collective behaviors to emerge from simple local rules, providing robustness and adaptability.

Planning and scheduling techniques are essential for optimal allocation of resources and real-time synchronization of multi-domain effects. They can constantly readjust execution plans according to changes in the operational situation and emerging opportunities.

Multicriteria decision-support tools make it possible to integrate logistical, tactical, strategic, and political dimensions into a holistic evaluation of options. These systems can dynamically weight different factors according to changes in operational context and command priorities.

3.2. Some Advantages for Land Operations

3.2.1. Accelerating Operational Tempo

The introduction of AI into land forces considerably accelerates the pace of operations. The ability to process information rapidly, make informed decisions, and optimize operational action and committed resources confers a decisive advantage on forces possessing these technologies. This compression of the decision cycle makes it possible to seize and maintain the initiative against a technologically less advanced adversary, who is always behind and unable to react to rapid changes in the tactical situation. Depending on the application and modes of action used, several effects can result from this acceleration: deterrence, discouragement, surprise, shock, and lightning tempo. This approach is well suited to French land forces and supports short, controlled engagements.

3.2.2. Reducing Human Losses

By automating the most dangerous tasks and improving early threat detection, AI can significantly limit soldiers' exposure and protect forces. Autonomous vehicles for intelligence gathering, reconnaissance, surveillance, logistics transport, demining robots, and shot-detection systems for snipers or artillery are all examples of AI-using technologies that preserve soldiers' lives. For the French Army, the maturity of autonomous vehicles is currently progressing along two axes:

- **Ground robotics:** personnel exposure is reduced through functions such as heavy-load carriage and persistent observation. AI requirements stress perception and automatic-navigation functions in unstructured, degraded, or complex environments.
- **Aerial drones:** endowed with strong mobility in the air domain, which is easier than the land domain, drones are nevertheless limited by computing power and communication discretion. AI providing simplified tactical autonomy helps protect operator personnel by limiting exchanges and performing elementary navigation tasks.

By extension, AI automates various functions of manned vehicles, such as close electronic warfare, self-protection, navigation, and counter-UAS operations. AI also provides embedded systems with very rapid reaction superiority and multi-sensor detection capability. In these systems, it is essential to return semantic information to humans–operators, users, and the chain of command–even when low-level inference is purely numerical.

3.2.3. Optimizing Scarce and Valuable Resources

Faced with budget constraints and often limited personnel, AI makes it possible to optimize the use of available physical resources. This includes optimizing capability potential, maximizing effects, planning routes according to different metrics–fuel, safety, speed, and others–prioritizing targets according to strategic importance, and optimizing equipment maintenance. Such AI systems have generally relied on constraint solvers and help maximize the operational effectiveness of land forces. In addition, systematic optimization of network resources and use of active sensors make possible the limitation of electronic warfare exposure.

3.2.4. Adapting the Tactical Level to Complex and Hostile Environments

Modern conflicts often take place in densely populated urban environments or in difficult, unstructured, or degraded terrain. Statistical AI helps land forces adapt to these complex contexts by providing tools for rapid terrain analysis, real-time mapping, and identification of potential threats. It enables systematic, rigorous, and objective identification of terrain changes, such as deployments, storage areas, transit routes, or fortifications. In addition, probabilistic AI facilitates information fusion and enables very rapid development of hypotheses in the face of simultaneous threats from different domains–space, air-land, cyber, and civilian. To enable deliberative decision-making, whether human or automated, these hypotheses must involve semantic predicates and symbolic variables.

3.2.5. Increasing Subsidiarity, Limiting Isolation, and Ensuring Resilience

Because perception-decision-action loops are accelerating sharply–owing to the massive use of fast drones, electronic-warfare techniques, and large-scale artillery–command must empower the chain in contact as effectively as possible by leaving lower echelons the capacity for initiative. In a heavily disrupted communications environment, the risk of isolation is high, reducing classical decision-making capabilities. Based on maintenance of the local situation, the use of AI assistance or automated vehicles helps preserve a resilient combat capability while respecting engagement rules and procedures, even under fragmented communications. In such cases, AI compensates for the absence of information superiority resulting from conventional tactical-communications schemes. Here again, AI must integrate symbolic knowledge of the tactical situation and account for command intent.

3.3. Data and Knowledge Management

Data and knowledge management is crucial to achieving the required reliability and performance. AI tools in defense involve software building blocks based on data and knowledge, which must be of high quality and available in sufficient quantity to design high-performing systems. Poor management of data, as well as knowledge, can compromise user trust and the effectiveness of tools.

Any policy for managing relevant data and knowledge will have to address several major challenges. Three major challenges stand out and require particular attention:

- the need to correctly represent ground truth, both in the data and in algorithmic results;
- the use of lessons learned to improve the quality and relevance of computational results or reasoning mechanisms; and
- the design of models from imperfect but realistic observations in order to guarantee robustness and reliability.

The data and domain knowledge underpinning this design, as well as model validation and updating, must be managed in a way that accounts for a series of imperatives. The following recommendations aim to guarantee appropriate data collection, annotation, and use; capture and representation of knowledge; and exploration of modes of collaboration between industry and the armed forces on this issue.

3.3.1. Data

Approaches to designing and updating data-driven AI models rely on machine-learning techniques that require high-quality data in sufficient quantity. The main data-quality issues are therefore availability and suitability of data, volume and variety of data, and freshness and security. Data must also be available, relevant, and adapted to operational needs, with appropriate structuring and standards. It is therefore essential to define a strategy suited to each use case, in collaboration between AI experts and operational personnel, in order to facilitate AI creation, prevent bias, and ensure compliance with operational requirements.

Datasets must be sufficiently large and varied to guarantee model robustness. It is therefore recommended to support initiatives for recording and annotating realistic datasets and to implement appropriate strategies, such as support for simulated-data generation, while ensuring data security and freshness. It then becomes possible to control learning biases and make models more robust. The creation of shared taxonomies and secure data sharing among different actors–state actors at minimum and potentially industrial actors within the DITB–are also essential to developing AI tools adapted to different professions and domains. These measures will make it possible to fully leverage AI for land operations.

Finally, models may require updates with recent data while guaranteeing the physical and cyber security of the data. It is therefore crucial to establish secure means for making data available and to define model-management rules to avoid any risk of regression, particularly for retraining conducted on the fly during operations.

3.3.1.1. Representing Ground Reality in Model-Training Data

AI algorithms must address one or more business or mission needs of their users. They must therefore be designed from the perspective of this use. Communication between algorithms and operators requires a shared and clearly defined language. For models integrating machine learning, the syntax of this language is formed by the taxonomies on which the models are trained and tested. These taxonomies must be defined collaboratively by AI experts and operational personnel in order to prevent data bias and ensure compliance with operational requirements. They must also be adapted to the different use cases considered, taking into account the specificities of land operations. They must be regularly updated to reflect changes in operational needs and technologies. They must also be shared among the different actors involved in the design and implementation of AI models to ensure coherence and interoperability among the tools developed.

Training data must be representative of ground reality in order to ensure model performance. They must therefore be: (1) collected, at least in large part, under real operational conditions to reflect the conditions in which models will be used; (2) annotated precisely and consistently by annotators trained in the domain to ensure the quality of trained models; (3) validated by domain experts to ensure the relevance of training data and the quality of annotations; (4) stored and managed securely to guarantee confidentiality and integrity; and (5) accessible to the different actors involved in model design and implementation so as to enable machine learning from these data and ensure tool coherence and interoperability.

Training data must also be sufficiently voluminous to guarantee the robustness of AI models, which need enough examples. They must therefore be collected in large quantities while ensuring diversity and representativeness; collected continuously when necessary to ensure freshness and relevance; and collected and annotated semi-automatically to ensure speed and effectiveness of the loop.

Finally, training data must be sufficiently varied to guarantee the robustness of AI models under the conditions in which they will be required to operate. They must therefore be collected across varied conditions and contexts in order to ensure representativeness and maximize coverage.

In practical and general terms, it is important that data be collected and annotated collaboratively to ensure annotation quality; stored and managed centrally to ensure accessibility and security; and shared among actors involved in the design and implementation of AI models in order to mutualize development costs as effectively as possible and make AI solutions as affordable as possible for the forces.

3.3.1.2. Use of Synthetic Data

Real observations under operational conditions are scarce, and data are often lacking both in quantity and quality. Alternative approaches can be considered to compensate for the lack of data, relying for example on synthetic data and/or frugal models, meaning models that need less data. Where appropriate, it is recommended to fund projects aimed at producing AI models trained on synthetic databases, to initiate data-sharing agreements among allied nations, and to invest in data-frugal methods.

Synthetic data must be: (1) generated realistically in order to ensure the relevance of simulated data for representing reality; (2) validated by domain experts to ensure the quality of trained models and their suitability for use cases; (3) subject to examination of their contribution to the predictive capabilities of AI models for the operational use cases considered; (4) stored and managed securely to guarantee confidentiality and integrity; and (5) accessible to the different actors involved in model design and implementation in order to promote coherence and interoperability among tools.

3.3.1.3. Use of Lessons Learned

Lessons learned are used to improve the quality and relevance of training data. They must therefore be collected systematically in order to monitor the relevance and reliability of predictions and improve AI-model performance; analyzed and validated by domain experts to ensure the relevance of their contribution to models for the use cases considered; stored securely and managed in successive versions so that model evolutions arising from them remain traceable; and shared among the different actors involved in the design and implementation of AI models to ensure coherence among the tools developed.

3.3.1.4. Design and Training of AI Models from Imperfect Observations

An inappropriate choice of training and/or test data, based on an insufficient understanding of rules, domain knowledge, and the intended use of AI building blocks, can compromise model performance. Insufficient data to train high-performing models or guarantee the performance of AI components trained on them, and knowledge that imperfectly or not at all covers the operational domain of the AI components, are challenges that must be addressed. Training data must therefore be adapted to operational needs and present all the appropriate properties mentioned above regarding collection, annotation, annotation control, dataset validation, storage, and sharing.

As for AI models themselves, they must be designed and trained to guarantee robustness and reliability. This means that they should be designed collaboratively by AI experts and operational personnel to ensure relevance and reliability. They must then be trained on high-quality data in sufficient quantity, validated by domain experts, and stored and managed securely to guarantee confidentiality and integrity.

They must also be updated regularly in order to ensure suitability for a continuously evolving domain need, because the context and threats are constantly changing. AI models must be updated according to field lessons learned, changes in operational needs, and technological evolutions. Such updates must be conducted securely to guarantee model confidentiality and integrity.

3.3.2. Knowledge

Symbolic AI, based on explicit rules rather than statistical predictions, is currently one of the most promising and controversial paradigms in the evolution of military technologies. Unlike statistical data-analysis approaches that currently dominate the AI landscape through neural networks and machine-learning algorithms more broadly, symbolic AI relies on the explicit manipulation of symbols and logical rules to represent and process knowledge. It differs fundamentally from other AI approaches through its method of representing and processing information. Whereas connectionist systems use distributed and numerical representations, symbolic AI uses explicit data structures in which each information element is represented by discrete symbols manipulated according to formal logical rules. This approach enables direct representation of expert knowledge as conditional rules, ontologies, and structured knowledge bases, organized into several broad families according to their approaches and methods.

Rule-based systems: This family uses explicit “if-then” rules to represent knowledge or human expertise as logical predicates.

Logic and automated reasoning: This branch relies on formal logic–propositional, predicate, and modal logics. It includes theorem provers, logical-resolution systems, and inference engines. The objective is to deduce new knowledge from established facts.

Knowledge representation: These systems focus on structuring and organizing knowledge: ontologies, semantic networks, frames, and conceptual graphs. They seek to capture the structure of knowledge domains in a formal and exploitable manner.

Automated planning: This family addresses the generation of action sequences to achieve objectives. Classical planning systems, such as STRIPS and PDDL, analyze states, actions, and goals in order to construct action plans.

Constraint programming (CSP): This family models problems as a set of variables, domains, and constraints to be satisfied. It includes constraint satisfaction, optimization under constraints, temporal constraints, and constraint propagation:

- constraint satisfaction: searching for solutions that satisfy all constraints, such as placement or scheduling;
- constrained optimization: maximizing or minimizing an objective function while respecting constraints;
- temporal constraints: managing temporal relations between events; and
- constraint propagation: techniques for reducing the search space.

CSP solvers use algorithms such as arc consistency, intelligent backtracking, or hybrid methods.

Fuzzy logic and approximate reasoning: This branch manages uncertainty and imprecision symbolically:

- fuzzy sets: membership of an element in a set is gradual rather than binary;
- fuzzy rules: “if X is approximately A, then Y is rather B”;
- fuzzy inference: deduction mechanisms with degrees of truth;
- fuzzy-control systems for industrial and robotic applications; and
- multivalued logics extending beyond classical true and false.

One could also add symbolic probabilistic reasoning–Bayesian networks with discrete variables and probabilistic logic–which combines symbolic aspects with uncertainty management.

In the specific context of land operations, symbolic AI has characteristics particularly suited to the complex challenges faced by modern armed forces. The highly structured nature of its reasoning mechanisms enables explicit modeling of military doctrines, rules of engagement, and standard operating procedures, thereby providing decision transparency that is often absent from connectionist AI systems. This transparency becomes critical when automated systems can directly or indirectly influence tactical and strategic decisions on the ground. Rules of engagement, for example, can be directly encoded as logical predicates, allowing the system to reason explicitly about the conditions authorizing the use of force. Similarly, movement and maneuver doctrines can be formalized as production rules guiding tactical decisions in real time.

The typical architecture of a military symbolic-AI system comprises several interconnected components: a knowledge base containing facts and rules relevant to the application domain; an inference engine capable of deriving new conclusions from existing knowledge; a knowledge-acquisition interface enabling the integration of new doctrinal elements; and an explanation system capable of justifying decisions taken. This last component is particularly crucial in the military context, where the traceability of automated decisions is a fundamental legal and ethical requirement.

That said, non-symbolic AI systems–based on machine learning or reinforcement learning–can also incorporate constraints such as rules of engagement or doctrine. These may be taken into account either gradually, by applying a greater or lesser penalty to predictions that do not respect them, or absolutely, by detecting with a conventional external judge system that a constraint has been violated, thereby explicitly invalidating the AI response, for example by forbidding a maritime drone from leaving a defined perimeter.

3.3.2.1. Knowledge Representation

Knowledge-based systems, which constitute the most mature manifestation of symbolic AI, find natural applications in the military domain. They can encapsulate the expertise of experienced commanders as heuristic rules, enabling the dissemination and standardization of tactical best practices. However, building such knowledge bases requires a complex and costly elicitation process, involving close collaboration between military-domain experts and knowledge engineers.

The first major challenge is transforming human knowledge into symbolic representations that can be used by machines. Land-military expertise is often tacit, acquired through experience rather than through explicit rule learning. Extracting and formalizing this expertise as symbolic rules requires a long and costly process involving domain experts who may have difficulty articulating their know-how explicitly. Moreover, this expertise constantly evolves with new threats, technologies, and lessons learned, requiring continuous updating of knowledge bases.

The challenge of completeness is particularly critical. It is practically impossible to represent exhaustively all knowledge in an operational domain, especially when that domain evolves. Knowledge-based AI systems operate in closed worlds, where what is not explicitly represented is considered false or nonexistent. This limitation can lead to incorrect reasoning in the face of unanticipated situations. In addition, imperfect information presents particularly acute challenges in the military environment. Traditional symbolic-AI systems struggle to manage situations where available information is partial, contradictory, or doubtful. Although extensions such as

fuzzy logic, Bayesian networks, or probabilistic-reasoning systems can mitigate these limitations, their integration significantly complicates system architecture and may compromise transparency.

Finally, modeling bias is omnipresent in symbolic systems. Every representational choice reflects a particular vision of the world. Formalized knowledge inevitably carries the cognitive and cultural imprint of its sources. Moreover, translating natural language into logical representations entails interpretive choices.

Maintenance and Complexity

Knowledge bases often suffer from maintenance problems. Adding new knowledge–business rules or operational constraints–can create inconsistencies with existing knowledge.

Depending on the technology used, combinatorial explosion can limit scalability. Symbolic-AI systems can quickly become unmanageable when the knowledge base reaches a critical size, creating performance and coherence problems. In the military context, where situations may require simultaneous consideration of many factors–terrain, weather, enemy, friendly forces, mission, and available time–combinatorial complexity can quickly exceed real-time processing capabilities. However, CSP-type approaches can exhibit the opposite behavior.

3.3.3. AI Vulnerabilities

Adversarial attacks[7] are a threat to AI systems. In symbolic AI, an adversary with access to the rules and logical structure of the system can potentially design inputs specifically constructed to exploit weaknesses or limitations in symbolic reasoning. Such attacks can be particularly sophisticated because they exploit the system's own logic rather than traditional technical vulnerabilities. Neural networks also suffer from this type of attack–for example, a slightly modified road sign can cause an autonomous car to fail.

Poisoning the knowledge bases of symbolic AI systems is another critical attack vector. If an adversary manages to introduce false information or malicious rules into the knowledge base of a symbolic-AI system, they can subtly but significantly influence the decisions taken by the system. This form of attack is particularly insidious because it can remain undetected for long periods while gradually compromising operational effectiveness. The situation is similar to the poisoning of databases used to train machine-learning AI.

Dependence on communications infrastructure also exposes these systems to denial-of-service and interception attacks. Existing military symbolic-AI systems often require real-time information exchanges with multiple sources, including satellites, drones, ground sensors, and centralized databases. Disruption of these information flows can seriously compromise system performance or even make systems dangerous by forcing them to operate with obsolete or incomplete information. This attack vector is of course also relevant to non-symbolic AI using the same information channels.

Symbolic knowledge bases raise clearer issues concerning the protection of sensitive data. Even without direct access to the data, a knowledge-based system can infer sensitive information through reasoning. Apparently benign rules may reveal correlations allowing classified information to be inferred. In addition, the transparency of symbolic systems, often presented as an advantage, can in some cases become problematic, because the ability to explain may expose confidential information. Finally, the combination of several knowledge bases can create unanticipated inferences, revealing sensitive information by cross-correlation.

4. What Responses?

4.1. Hybridization

Today, the challenges induced by data-driven AI and knowledge-based AI take on a new dimension with hybrid AI and more particularly the emergence of neuro-symbolic approaches. They aim at combining machine learning with symbolic reasoning, taking advantages of each paradigm while mitigating their respective limitations. Future systems will probably combine the advantages of symbolic AI–transparency, explainability, doctrinal compliance–with those of connectionist approaches–learning capability, robustness to noise, and generalization. This convergence could overcome some current limitations while preserving the specific benefits of each approach.

[7] An adversarial attack consists in adding a small, imperceptible perturbation to an input in order to modify the output of a machine-learning model, such as altering the classification of an image.

Hybrid Approaches
A combination of several AI methods–for example symbolic and statistical AI–possibly with other mathematical or physical approaches, in order to leverage the advantages of each: the data-processing power of statistical AI; the transparency and traceability of symbolic reasoning; and knowledge of physics or mathematical properties such as rotational invariance. Example: an assisted command system that uses deep learning to analyze images and then applies formalized rules of engagement to propose a regulation compliant decision.

4.1.1. Physics-Informed Neural Networks

The new class of neural networks called Physics-Informed Neural Networks (PINNs) is based on hybridization with physical models. These neural networks are trained to solve supervised-learning tasks while respecting all the laws of physics described by differential equations, which limit the solution space admissible to the neural network during learning. The physical laws integrated into PINNs can be highly diverse, ranging from fluid mechanics, such as the Navier-Stokes equations, to electromagnetism, such as Maxwell's equations, and thermal phenomena, such as Fourier equations. More generally, any physical law expressed as differential equations can be captured by a PINN. Although this new technology has not yet reached technical maturity, with a low TRL, it will have important applications in coming years in many fields such as aerospace and defense, including through digital twins.

4.1.2. Geometry-Informed Neural Networks

Information geometry has become a very popular tool in AI, particularly among major technology companies, which use the gradient associated with the Fisher metric to account for the geometric structure of multilayer-network parameter space. Early proofs of concept for Geometric-Informed Neural Networks (GINNs) have been implemented for functions such as Automatic Target Detection and Recognition on micro-Doppler or kinematic target signatures, or image recognition from 360-degree fisheye cameras.

4.1.3. Deep Morphological Networks

Convolutional neural networks (CNNs) have proven effective for object classification in images. In some cases, however, CNNs perform poorly, particularly in terms of geometric interpretation. To overcome this problem, non-linear operations such as morphological operations can be used. Mathematical morphology analyzes an image through a structuring element whose topology or geometry is controlled, using elementary set operations such as erosion, dilation, opening, closing, and morphological gradient. This knowledge-based AI technique, whose knowledge is captured by the structuring element, transforms the image progressively to bring out elements of interest. Its hybridization with neural networks was applied to shape recognition as early as 1991, but in recent years a new neural architecture has appeared: the Deep Morphological Network, or DeepMorphNet. In such networks, convolutions are replaced by morphological filters capable of performing non-linear operations while learning the characteristics underlying the structuring elements. DeepMorphNets can therefore model spatial relations of an object of interest or learn the structure of an image.

4.1.4. Fuzzy Inference

More unusual combinations of AI can also be implemented. Fuzzy symbolic systems, for example, integrate fuzzy logic and knowledge-based systems. These systems extend rule-based systems by adding the possibility of representing fuzzy rules and manipulating them through fuzzy-inference mechanisms. Neuro-symbolic-genetic systems are composed of a genetic algorithm responsible for acquiring knowledge from data–learning–and a symbolic module responsible for symbolic inference–reasoning. In this movement, the Alpha module from the U.S. start-up Psibernetix, acquired by Thales in 2019 and known for defeating experienced fighter pilots in air-combat simulations, is based on a combination of fuzzy logic, decision trees, and genetic algorithms, thereby providing a degree of resilience to noise, environmental uncertainty, and various contingencies.

4.1.5. The Benefits of Hybrid AI

Hybrid artificial intelligence is expanding impressively, driven by powerful promises such as robustness, explainability, frugality, and use in collaborative decision-making.

While each form of superiority benefits from specific AI technologies, decisive advantage arises from integrating them into architectures adapted to specific needs–often hybrid architectures–covering the entire informational, decision, and operational spectrum:

- By combining different AI approaches across tactical, operational, and strategic echelons, hierarchical cognitive systems enable vertical decision coherence while preserving adaptive autonomy at the local level.

- By integrating different specialized technologies to coordinate effects across physical and informational domains, these AI systems allow precise synchronization of actions in land, air, electromagnetic, and information dimensions.

This holistic approach to the hybridization of AI technologies–symbolic, connectionist, and generative–is the most promising path for developing sustainable military superiority in contemporary and future conflicts, where informational, decision, and operational dimensions are inseparably related.

4.2. Trustworthy AI and Performance Assurance

Under the new European regulation, the AI Act, an AI-based system is trustworthy if it meets six high-level requirements: robustness, effectiveness, reliability–including safety and security–usability, human-system interaction–including transparency, explainability, and interoperability–and human control, including ethical issues.

Trustworthy AI
AI designed to be reliable, explainable, robust, ethical, and controllable by its human operators. In a military context, this includes traceability of decisions, resistance to adversarial manipulation, and compliance with rules of engagement. Example: a decision-support system that indicates not only its recommendation but also its confidence level, the data used, and the limits of its reasoning.

The characteristics of trustworthy AI can therefore be defined as follows:

- **Robustness:** the system's ability to perform the intended function in the presence of abnormal or unknown inputs.
- **Effectiveness/Correctness:** the ability to fulfill the functions required to meet the requirements.
- **Reliability:** the ability to provide a service that can be justifiably relied upon. This property concerns not only the system itself but also the other actors and processes involved throughout the AI life cycle–engineers, operators, certification authorities, insurers, and others.
- **Usability:** the extent to which the system can be used to achieve objectives with efficiency and satisfaction in a specific context of utilization.
- **Human-system interaction:** the ability of individuals to interact with AI-based systems, understand them, and control them, ensuring that these technologies are transparent, explainable, and aligned with human intentions.
- **Human oversight:** the assessment and guidance of AI-based systems to ensure that their operation respects legal frameworks, fundamental rights, and general benevolence.

These requirements take specific forms for critical systems, particularly in defense. Safety may, for example, require formal validation and even certification for certain application contexts. Real-time self-explanation may be necessary for the acceptability of certain AI-based critical systems and may require advanced human-machine dialogue. Cybersecurity has a complex bilateral relationship with AI. Some characteristics of learning systems make them vulnerable to cyberattacks and deception, including AI-driven deception. Conversely, some AI algorithms can identify irregularities or anomalies and thereby help prevent cyberattacks. All these issues raise concrete technological challenges.

To guarantee the algorithmic design of trustworthy AI, integrating AI paradigms as well as quality, safety, and cyber-security dimensions requires demonstrating that algorithms are correct. It is necessary to verify compliance between specifications and behavior: the gap between what the system is supposed to do and what it actually does. Some symbolic-AI approaches, such as constraint programming, provide this correctness property by construction. For connectionist AI, which is stochastic in nature, this demonstration is generally performed through test campaigns. Moreover, because the robustness of an AI-based system characterizes its ability to provide correct responses in the face of unknown situations or malicious acts, it is more difficult to qualify than accuracy. A system that is not accurate cannot be robust. But a system that is accurate may still fail to be robust. This is the case for a learning-based system that has memorized its training data and then makes errors in future decisions based on new data; this phenomenon is known as overfitting. AI also remains vulnerable and, if care is not taken, particularly sensitive to adversarial attacks, which exploit the functioning of underlying algorithms to generate low-amplitude perturbations in analyzed data and force the AI to return an incorrect result. Fortunately, the existence of adversarial attacks implies the existence of defenses. Many defenses have been proposed in recent years by the scientific community, although some are later refuted by new attacks that make them obsolete. Some hybrid-AI approaches are therefore more robust. Examples include PINNs, which are robust to sensor noise such as vibration, and GINNs, which are robust to geometric deformations such as distortion induced by fisheye cameras.

It is also necessary to prove that critical systems are controllable, that is, well-founded or consistent, if one can prove that they do only what is expected of them. Questions concerning robustness and consistency are beginning to be addressed by work on formal proofs. These aim at providing "a priori" guarantees on the safe

operation of a program, unlike validation methodologies based on direct experimentation, that try to provide a posteriori guarantees. This is why the combination of symbolic and connectionist AI appears very promising.

The famous “black box” of AI is a major concern for future developments. For critical systems, it is necessary to fully understand how and why an algorithm makes a decision. The Villani report, *For a Meaningful Artificial Intelligence*, emphasizes the importance of this issue, breaking the challenge into three axes: producing more explainable models, producing more intelligible user interfaces, and improving understanding of the underlying cognitive mechanisms. In practice, one can envisage building hybrid systems that intelligently mix symbolic AI, such as knowledge-based systems, with learning-based approaches.

4.3. Integration and Embeddability of AI in Operational Systems

Integrating AI-based software solutions into embedded systems is a complex subject at the boundaries of several domains.

4.3.1. Development Methodology

First, integration into an embedded system must be considered as early as possible during design phases. It is necessary to begin by managing the possible gap between the design environment and the environment in which the function will be deployed. In practice, this covers two themes: data and hardware.

The data used to build the training base for the AI function, or more generally the data used during design, may differ from the data encountered once the system is deployed.

During the early stages of design, it is therefore necessary to choose the source used to build the training base and to choose which data will be used to qualify the system. This means, among other things, determining the proportion of synthetic and real data, defining the content of acquisition campaigns to be conducted, and listing the field trials needed for qualification. The purpose is of course to ensure correct system operation once deployed.

On the hardware side, it is important to note that the development strategy has a strong impact on the lessons and qualification arguments that can be produced. Producing a demonstrator for risk-reduction purposes is very different from producing a product prototype.

More practically, porting AI functions onto embedded computers often involves compression steps, especially for real-time applications with complex computations. The functional impact of this compression step must be accounted for in the development process, particularly when qualifying the system, which must assess the behavior of the function as it will actually execute on the final computer in the deployed system.

4.3.2. Adaptation to Embedded-Computing Constraints

Embedded computing is limited, among other factors, by available computing power. The use of AI models, which are often computationally intensive, must therefore be justified. Matching algorithmic complexity to the actual need frees valuable computing resources. This can be achieved, for example, by selecting models tailored to the task or limiting the employment domain covered by the function.

It must not be forgotten that, in most cases, the AI function will not be alone on the computer. Most often, such AI functions are integrated into conventional software chains that ensure overall coherence or provide other functions executed in parallel. This issue also enters into the selection of the computer, which must be able to manage associated constraints: frequency, latency, reproducibility, and so forth.

Computing resources available alongside the AI function can also be mobilized to monitor that function. This may consist, for example, in monitoring function outputs or monitoring inputs to verify that they are compatible with the qualification domain, the Operational Design Domain (ODD).

4.3.3. Cybersecurity and Information-System Security Issues

Securing software content changes scale with AI, because the volumes of data to protect can increase compared with traditional algorithms. The parameters of AI models are sometimes numerous and voluminous and must be protected; otherwise they may reveal clues about how the function operates or about its performance. Such protection sometimes requires compromises on execution performance.

4.3.4. Ability to Update AI Functions in Embedded Contexts

Traditionally, the AI function is initially trained before being integrated into the system, qualified, and deployed. This static view, already an industrial challenge, is now evolving with the emergence of a need to update the function and therefore to perform new learning.

This new learning will have to occur fairly frequently, and on or near the deployment field. This need involves several issues:

- availability of data to refine function operation and guarantee the benefit of the update;

- access to computing resources for additional learning and updating;
- depending on context and users, partial or full automation of the field update, qualification, and non-regression verification process to guarantee deployment; and
- access to the systems and ability to load new parameters into them.

4.4. Human Appropriation of AI

The use of AI in military operations raises important ethical questions, particularly concerning the autonomy of weapon systems. The fundamental principle of maintaining the human in the decision loop for the use of lethal force remains a major concern. Land forces must develop clear doctrines and appropriate legal frameworks to regulate the use of these new technologies.

Doctrine alone, however, will not suffice. Forces and industry must work together to design optimized systems allowing the operator to understand the situation and make an informed decision. These systems must include elements guaranteeing trust in system operation and predictions, as well as well-designed human-machine interfaces. To maintain situation understanding, these HMIs must provide the right information, at the right level of synthesis, at the right time. This maxim, already difficult to master for a subsystem, becomes even more demanding when managing multiple systems, or even swarms or packs.

4.4.1. Use and Human Factors

Human appropriation of AI rests on an alliance among performance, traceability, and explainability–elements that are crucial to establishing trust between humans and machines. This trust is indispensable for full exploitation of the assistance provided by AI. By 2035, performance and traceability levels appear attainable, but explainability remains a major challenge, especially in defense, where weapon systems require clear justifications for their results. Integrating AI into robotic systems transforms the role of the robot from a mere tool into a true teammate. Today, reluctance to use robots integrating automated functions persists, particularly because of fears of losing human-to-human communication and of facing uncontrollable robots making decisions autonomously.

4.4.2. Teleoperation and Supervision

Teleoperation consists in performing operations remotely, without the physical presence of the operator at the time and/or place of operations. This distinction is crucial because it involves complex human mechanisms such as remote perception and mental representation. Supervision, for its part, consists in monitoring and diagnosing the operation of a system by collecting data that provide indications about its state. Operators must not only understand how to interact with automated systems but also how such systems can influence their perception and cognition. Remote perception and mental representation, for example, play a crucial role in how humans perceive and interact with virtual or distant environments.

4.4.3. Cognitive Load and Human Factors

Human factors such as cognitive load, perception, and vigilance are key elements in the appropriation of AI. Operators must be able to understand and manage the information provided by automated systems, which requires appropriate training and experience. Communication constraints and performance requirements for critical tasks are also important factors to consider. A visuo-vestibular conflict, for example, can cause cybersickness, spatial disorientation, and loss of spatial reference, complicating AI appropriation. The same situation is found when using mobile automated systems such as ground robots or aerial drones, with or without AI assistance.

This evolution requires a clear organization of tasks and roles within the human-robot system while accounting for the operator's cognitive capabilities. The coexistence of different environmental reference frames and the operator's ability to supervise several platforms while maintaining a global view are cognitive and organizational challenges to be met. To do so, it is crucial to delegate some tasks to the system, allowing the operator to focus on specific tasks while retaining the ability to take back control if necessary. Operator involvement in supervision tasks is also essential, driven by the need to understand system decisions, know system state, and remain in control. These mechanisms of trust and technology acceptance are often reinforced by change-management plans, facilitating the integration of AI into work habits.

4.4.4. Explainability and Trust

AI must enable better understanding of the situation and greater responsiveness in decision-making. For this relationship to be effective, however, AI must be explainable. Deep-learning systems, although effective, must be able to justify their decisions in order to earn operators' trust. Ongoing work on AI explainability is essential to overcoming this challenge and ensuring the successful integration of AI into military systems. Given the

potentially serious effects of weapon systems, work aimed at making artificial-intelligence systems capable of justifying their own results must be viewed as a prerequisite for exploiting AI technologies.

4.4.5. Organizational Considerations

Integrating AI into armed forces requires a profound transformation of military culture and of the skills required. Land forces must not only train their personnel to use these new technologies but also help them understand their limitations and maintain critical judgment regarding their recommendations. This cultural change can sometimes represent a greater challenge than technological development itself.

Recent work has shown that operational personnel have varied expectations regarding automation and AI. Some prefer direct supervision of automated-system actions, while others are more inclined to delegate specific tasks to AI. Reluctance and fears also persist, particularly due to lack of confidence in the proposed technology and difficulties in referring to and accounting for observation data across distinct terrain sectors, which complicates supervision and teleoperation. These differences of opinion highlight the need for in-depth training and adaptation of organizational processes to integrate AI effectively into military operations.

Furthermore, concerns exist regarding operational responsiveness in the presence of automated systems, because the high cadence of tasks increases perceived stress and cognitive load across the crew. These concerns are linked to the perception that current technology cannot react as quickly as human operators during intense engagement phases.

Finally, it remains necessary to have a particularly solid technical base to respond to the contingencies and malfunctions of automated systems. This implies continuous training and adaptation of operator skills. In short, integrating AI into armed forces requires a holistic approach that accounts not only for technological aspects but also for cultural, organizational, and human dimensions.

4.5. Sovereignty

AI sovereignty refers to the ability of a country or region to retain control over its infrastructure, data, production and maintenance capacity, and AI-related decision-making processes. For AI suppliers, this raises several important questions that they must address.

Data localization and storage: Governments increasingly require sensitive data to remain within national borders. AI companies must therefore build local data centers, establish in-country processing capabilities, and ensure compliance with data-residency laws. This can significantly increase infrastructure costs and operational complexity.

Access to data: Data-driven AI, especially machine-learning-based AI, requires large volumes of data for training. Access may be blocked by factors such as classification, export control, trade secrets, or regulation. The entry cost for a market actor therefore depends on its ability to access the data necessary to develop its product or service. Simulation environments seek to bypass this difficulty, but deploying a system trained only on synthetic data raises several questions, including reliability. A balance between real and synthetic data remains necessary.

Access to infrastructure for effective data management: Upstream, for training, and downstream, for operational exploitation, access to hardware and software infrastructure is indispensable. Here again, the largest actors are currently American, including the major U.S. technology companies. In addition, the software technology stack covering the MLOps/AIOps chain, also dominated by U.S. actors, mixes open-source and proprietary technologies whose conditions of use may change rapidly.[8]

Alternatives to U.S. clouds exist. Nevertheless, if the entire industrial fabric had to use them rapidly, scaling issues would arise.

Regulatory compliance and standards: Different jurisdictions are developing their own AI regulations, such as the European AI Act, Chinese AI regulations, or new U.S. regulatory frameworks. Companies must adapt to compliance requirements, security standards, and approval processes that vary by market, often requiring different versions or features for each region. Industrial efforts in standardization committees can help harmonize rules across different zones.

Technology transfer and export controls: Many countries impose restrictions on exports of AI technologies, especially advanced models or those with dual-use applications. Companies face challenges concerning which

[8] Open-source licenses for YOLO vary from permissive licenses such as Apache 2.0 and MIT to restrictive licenses such as GPL-3.0 and AGPL-3.0, and even commercial licenses. Apache 2.0 is the most flexible for commercial or proprietary use, while GPL-3.0 and AGPL-3.0 promote open collaboration with obligations to share modified code. AGPL-3.0, used for YOLOv8, is strict and requires any modification or SaaS/cloud deployment also to be open source. GPL-3.0, used for YOLOv3, YOLOv5, YOLOv6, and YOLOv7, requires any derivative work also to be open source under the same license. Apache 2.0, used by YOLOX, PP-YOLO, YOLO-NAS, and others, allows free use, modification, and distribution, including in proprietary projects, without the obligation to open source modified code. The MIT license is also highly permissive.

technologies they may share, where they may deploy certain capabilities, and how to handle government requests for access to technology or source code.

Local-partnership requirements: Some governments require foreign AI companies to partner with domestic firms or create local subsidiaries in order to operate in their markets. This may involve sharing intellectual property, training local talent, or relinquishing some control over operations.

Transparency and explainability requirements: Governments may require AI companies to provide detailed explanations of how their models work, their training-data sources, training processes, or decision-making processes. This may conflict with commercial interests and competitive advantages.

Hardware supply: The civilian ecosystem facilitates rapid porting of AI functions to certain embedded computers. This is especially true for GPUs, which make it easy to produce embedded demonstrators. However, the tools used for porting, such as optimization or compression libraries, are often opaque and tied to the computers used. This limits the ability to design embedded algorithms with full control.

In response, ongoing work is enabling the emergence of sovereign tools such as AIDGE, led by the French Alternative Energies and Atomic Energy Commission, which guarantees transparency and control in design.

As for the choice of the computer itself, many criteria come into play: ITAR regulations, cost, size, power consumption, computing power, environmental resistance, ease of use for algorithmic and software development, cybersecurity vulnerability, information-system-security constraints, production and supply, maintenance, scalability, and more.

Producing and deploying AI requires technical means, especially hardware. These means are currently supplied mainly by American actors, themselves still heavily dependent on semiconductor supply from Asia, especially Taiwan and South Korea. The industry risks being prevented from acting in the event of supply disruption, or being downgraded in price competitiveness if financial access conditions deteriorate, for example during a trade war. French initiatives, such as those led by CEA, and European initiatives, such as the European Commission's Chips plan, seek to mitigate these weaknesses. But it must be acknowledged that there is not yet a scalable fallback solution to our current dependence. Supply disruptions during COVID demonstrated the full weakness of this dependence for industrial production, especially in the automotive sector.

Choice of models: To produce deep-learning AI, model architectures are required. These architectures are largely made available free of charge and open source. Nevertheless, several major actors are beginning to monetize access to their architectures–Ultralytics in computer vision, for example–or may no longer publish them, as with companies producing large language models. Access to high-performance architectures is a necessary condition for obtaining quality models.

Today, the great majority of architectures are invented in American companies or laboratories. Here again, tighter access could delay industrial development.

Economic and strategic dependencies: Countries fear becoming too dependent on foreign AI suppliers for critical infrastructure or decision-making. AI companies must address concerns about continuity of service, price control, and potential geopolitical risks that could affect service availability. Such sovereignty concerns often require AI suppliers to make significant investments in local infrastructure, adapt their business models to different markets, and balance global efficiency with regional-autonomy requirements.

As a synthesis, industrial AI sovereignty aims to preserve the ability to act–freedom of action–for actors in a given sector. Sovereignty of an industry, according to this definition, can therefore be broken down across several axes that tend to cover the reality of its value chain or manufacturing process.

5. Recommendations

5.1. Six Years Later

The military and technological contexts have changed radically since GICAT's previous report on AI for air-land systems, published in 2020, imposing a profound transformation of approaches to integrating artificial intelligence into defense systems. At the time, the objective was already clear: convert data and knowledge into actionable information. But the approach remained analytical, segmented, and exhaustive, targeting sectors and use cases where AI promised immediate returns, particularly in command and control and intelligence. Yet despite progress, it must be acknowledged that the initial ambition has been only partially achieved, even though the emergence of generative and multimodal AI now offers new perspectives for fusing heterogeneous data and accelerating operational adoption.

In the meantime, France has established AMIAD, the Defense Artificial Intelligence Agency, marking a desire to structure this transition. But the real paradigm shift lies in the transition from an analytical logic to a holistic and systemic approach. The issue is no longer merely to identify technological niches, but to rethink AI as an

integrated ecosystem in which automation, delegation to the machine, and system resilience become central issues. Lessons from the Ukrainian theater have confirmed this necessity: military innovation is now driven by civilian actors, often more agile, while armed forces must cope with accelerated development cycles and unprecedented technological competition.

The points of vigilance identified six years ago remain crucial: sovereignty over data and infrastructure, volume and quality of learning corpora, hybridization of AI techniques, and adaptation to degraded environments. Yet the lead of the civilian sector has widened, and the war in Ukraine has revealed how non-state or private actors can disrupt strategic balances by rapidly transposing dual-use innovations. These six years have demonstrated one reality: without strong political will and systemic coordination, France will not be able to compete with competitors that integrate AI at a sustained pace and without the ethical or regulatory constraints that weigh on democracies. The issue is no longer solely technological, but strategic and organizational: moving from a logic of silos to a global vision in which AI becomes a capability multiplier rather than a mere tool.

5.2. Today's Major Challenges

5.2.1. Data and Knowledge: The Winning Combination

In the demanding context of land-forces military operations, hybrid AI, combining data-driven approaches with knowledge-based approaches, is particularly relevant to addressing the specific challenges linked to data availability, system robustness, and the integration of domain knowledge. It is often difficult to obtain fresh, complete, and representative data–particularly for tasks such as classification or automatic threat detection–because of the sensitive nature of such data or its real-time unavailability. It therefore becomes essential to exploit "opportunity data": older data, data from exercises, and less sensitive but still valuable data. Working only with data, however, is not enough. Domain knowledge–the operational, tactical, and strategic expertise specific to the forces–provides a unique decision advantage that is difficult for other actors to reproduce.

To maximize the benefits of AI, it is also crucial to integrate domain knowledge from the earliest phases of system design, in parallel with the development of technologies such as AI and sensors, so that models are guided by operational needs and constraints. This makes it possible to design systems that are more effective, better adapted to field contexts, and more explainable.

Technically, conventional AI approaches still present notable limits. Statistical and connectionist methods based on large quantities of examples often lack transparency and interpretability, which reduces the confidence of operational users. They are also insufficiently robust when faced with novel or adversarial situations and consume large quantities of data and energy–resources that may be scarce in deployed contexts. Conversely, symbolic AI, which uses explicit rules and formal reasoning, provides better explainability, but its robustness to uncertainty or unmodeled situations remains limited. Combining the strengths of connectionist approaches–neural networks and deep learning–with symbolic approaches, and integrating available domain, physical, and mathematical knowledge, increases model interpretability and operational robustness and facilitates validation for the approvals required for critical military systems.

Research programs such as DARPA's Assured Neuro Symbolic Learning and Reasoning (ANSR), launched in 2022, embody this approach by developing algorithms that intimately integrate symbolic reasoning and machine learning in order to produce reliable, safe, and trustworthy systems.

Among the most promising innovations, Physics-Informed Neural Networks hybridize neural networks and physical models. By constraining learning to respect fundamental laws of physics modeled through differential equations, PINNs reduce the space of possible solutions and increase reliability and explainability of results. This technology, applicable for example to fluid mechanics, electromagnetism, or thermal phenomena, opens significant prospects in defense, particularly through digital twins, which make it possible to simulate in real time and with high precision the complex behaviors of equipment in operational environments. Geometric-Informed Neural Networks use information geometry to better capture the structure of multilayer-network parameter spaces. This technique has already enabled proofs of concept in critical military applications such as automatic target detection and recognition based on micro-Doppler signatures, fine analysis of kinematic trajectories, and visual recognition from 360-degree fisheye cameras.

Intelligent exploitation of data, combined with the valorization of domain knowledge and rigorous integration of physical and mathematical laws, makes it possible to build AI systems suited to the requirements of robustness, interpretability, and trust that are indispensable to land forces. This approach opens the way to enhanced battlefield capabilities by strengthening decision-making, early threat detection, and efficient resource management in an increasingly complex and uncertain environment.

5.2.2. Cyber Vulnerabilities

Increased dependence on intelligent systems creates new vulnerabilities to cyberattacks. An adversary capable of compromising or disrupting AI algorithms could neutralize technological advantage or, worse, turn these

systems against their users. Land forces must therefore develop robust cybersecurity protocols and degraded-mode operating capabilities.

Among the most important threats are attacks aimed at modifying model behavior. Securing AI learning on sensitive data against information leaks and protecting copyright in a complex context are also major concerns. Pending regulation, professionals are seeking countermeasures that combine protection and performance.

To ensure traceability of a shared model, its creator can use watermarking techniques inspired by media-marking solutions. AI-model watermarking aims to protect intellectual property by integrating proof of origin into the architecture or behavior of the model. Watermarking techniques fall into two main categories: white-box and black-box. In a white-box setting, proof of ownership relies on a secret modification embedded in the model parameters or architecture. In a black-box setting, it takes the form of a secret modification embedded in model behavior.

Federated learning enables models to be trained directly on local devices, with updates aggregated into a global model without raw data leaving each device. This facilitates collaboration and reduces data-security risks. Federated learning, however, does not protect the final model from information leaks, especially if an attacker compromises the aggregation server. Homomorphic encryption and differential privacy techniques can be used to mitigate this risk.

An effective security strategy begins by identifying the assets to be protected, such as training data, the model, or input/output data. It continues with an assessment of potential threats and impacts according to the attacker's capabilities. In a white-box scenario, the attacker has access to the entire model; in a black-box scenario, the attacker can only query it. Corrective measures can then be implemented to improve model robustness or add defenses to the system.

Data poisoning is another major threat. An attacker with access to a model's data supply chain may seek to manipulate it in order to poison the model, thereby degrading performance. Poisoning can also be used to insert functionality unknown to the legitimate user, known as backdoors. One notable example is Microsoft's TAY chatbot, which had to be taken offline after malicious internet users made it produce hateful statements.

Generative AI also raises specific challenges. Large language models are a new target for cyber attackers. Prompt injection allows attackers to use carefully crafted inputs to manipulate the LLM into executing potentially malicious instructions. This can lead to manipulation of model responses or of any decision-making process that the model influences or controls. Jailbreaking attacks consist in asking AI to adopt a different identity to discuss illegal acts, hateful content, or disinformation. Deepfakes generated by chatbots can be used to deceive biometric systems, carry out social-engineering attacks, or conduct information warfare.

5.2.3. Interoperability and Standardization

Deploying AI is not simply a matter of installing and executing code in isolation. AI systems will be deployed within systems integrating sensors, effectors, users, data links, and non-AI components, all in interaction. The architecture of these systems must therefore account for how these components interact with one another. To promote interoperability, several stages can be distinguished in the life cycle of AI systems where standardized interfaces, or at minimum standardized descriptions, are needed:

- **Design:** the types and formats of data used by each component, the intended employment domain of the system, and the methods used to design it, such as the type of training in the case of learning.
- **Validation:** the performance achieved and known limitations of the system.
- **Integration:** hardware constraints required to guarantee performance.
- **Production:** the formats of data streams produced and any usage alerts.
- **Maintenance:** the formats of captured data and the corresponding employment domain, achieved performance, and detected incidents.

Standardizing these formats, interfaces, and descriptions facilitates use by actors across the value chain: operators, integrators, maintainers, suppliers, and deployers. Today, AI systems are sometimes designed in isolation from the rest of the system, and their description often remains dependent on the supplier, frequently from the civilian world, who may still define proprietary description formats. In some cases, the information provided does not allow full assessment of usage frameworks or system performance, because only part of the necessary metrics is provided. Standardization efforts are under way in the civilian world, for example in CEN-CENELEC JTC 21 and ISO/IEC JTC 1/SC 42, and such documents may serve as a basis for a military adaptation of standards within a framework such as NATO or the European Defence Agency.

In a context of joint and multinational operations, interoperability of AI systems is a crucial issue. Harmonizing data formats, communication protocols, and human-machine interfaces is necessary to guarantee the operational effectiveness of coalition forces.

5.2.4. Explainable AI

AI systems are often described as black boxes from which it is difficult to extract justifications for what they produce. This observation is broadly true for machine-learning AI and especially true for deep learning. Some more symbolic techniques, however, intrinsically carry a capacity to justify their outputs.

For a user, trust in a system is partly built through a relationship with that system in which the user can understand and anticipate its actions. ISO/IEC distinguishes two levels related to this understanding: explainability and interpretability.[9] In interpretability, the system provides low-level information used to analyze which inputs contribute to producing the outputs. This level is generally used by system designers to verify behavior and make corrections where needed, including for machine learning and especially deep learning. In explainability, the system can provide an analysis understandable by a user who is neither an AI expert nor a system designer–therefore close to the level of human reasoning and possibly in natural language.

To strengthen operator trust and maintain meaningful human control, future military AI systems will need to explain their reasoning and recommendations clearly and intelligibly. To do this, user expectations must be measured, as must users' level of training. Explainability must be adapted to the user's level of understanding and to the time available for the user to analyze it.

5.2.5. Rethinking the Engineering of AI-Based Systems

Many barriers still hinder AI adoption, particularly deployment in critical systems. These systems must, by construction, guarantee properties of reliability, maintainability, availability, safety, and security,[10] while also following principles of ethics and responsibility. For land operations, it is essential to deeply revisit the different engineering disciplines: algorithmic engineering, software engineering, systems engineering, data engineering, knowledge engineering, cybersecurity engineering, safety engineering, and cognitive engineering. AI integration often follows an experimental approach, requiring traditional working methods to be rethought. It is no longer enough to validate a point solution; instead, incremental integration and qualification must be favored so systems can be progressively adapted while guaranteeing robustness.

AI is also highly contextual by nature, requiring adaptation to changing operational contexts. This implies permanent dialogue among operational users, systems engineers, developers, and data scientists in order to identify genuinely relevant data and knowledge–for example from exercises–and to systematize their extraction, collection, qualification, and exploitation. In particular, equipment and systems must be equipped to support effective data collection and use.

In the specific defense context, a major issue is avoiding reverse engineering in the event of system compromise by an enemy. This strongly influences architectural and security choices around embedded AI.

Finally, while a proof of concept for an AI application is often easy to obtain, transitioning to a reliable industrial solution adapted to military constraints is complex. AI engineering must therefore integrate not only scaling but also continuous maintenance, regular updates, and guaranteed performance over time. This global transformation of engineering processes is an indispensable prerequisite to fully exploit AI's potential in critical systems, including for land operations.

[9] In ISO/IEC 22989:2022–AI: Overview of trustworthiness concepts–explainability is defined as the ability of an AI system to provide, in a human-understandable way, information about internal functioning that justifies a decision or behavior, while interpretability represents the degree to which a human can grasp cause-and-effect relationships among inputs, internal processes, and outputs in order to understand why a decision was made.

[10] RAMS: Reliability, Availability, Maintainability, Safety.

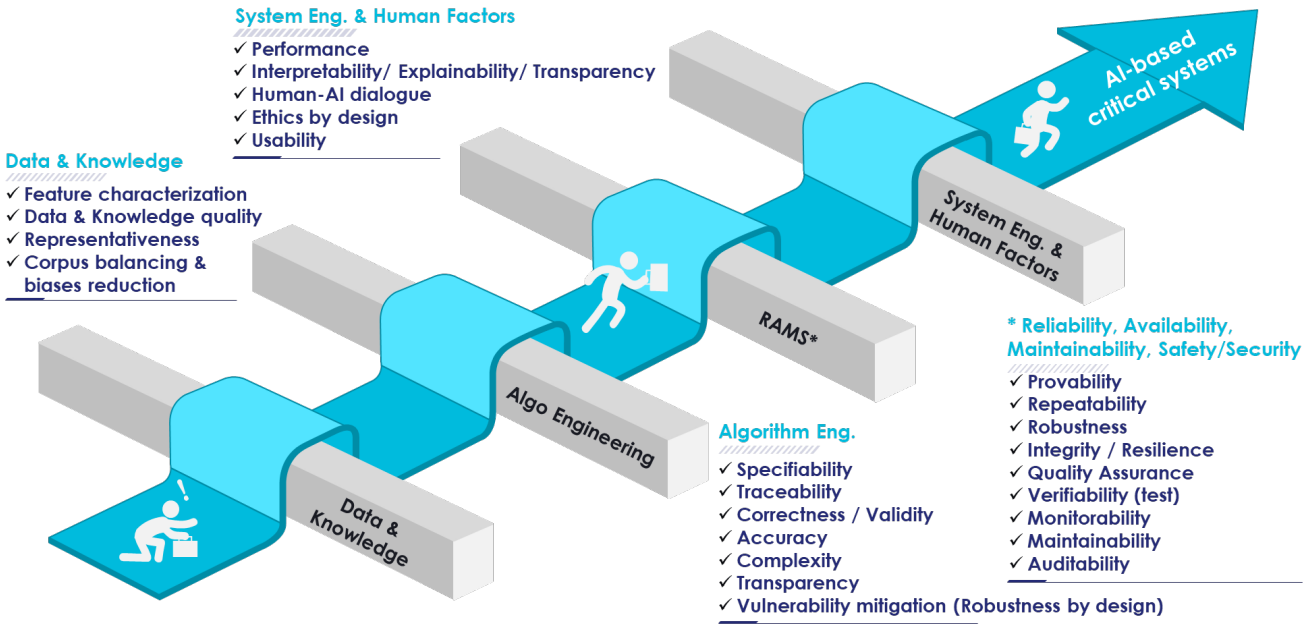

Figure 16: The introduction of AI creates new engineering challenges.

Thus, the development of an AI-based system must rely on well-founded development methods from design through deployment and qualification, creating new engineering challenges.

Engineering practices must therefore be revisited and enriched with methods and tools guaranteeing trust at every stage of the life cycle of such a system: (1) analysis of the operational domain (OD) with respect to the intended purpose; (2) specification of the Operational Design Domain (ODD) and its derivation for data and knowledge management; (3) design of algorithms and architecture; (4) characterization, verification, and validation; (5) deployment, particularly on an embedded architecture; (6) qualification and certification; and (7) sustainment of operational condition and cybersecurity.

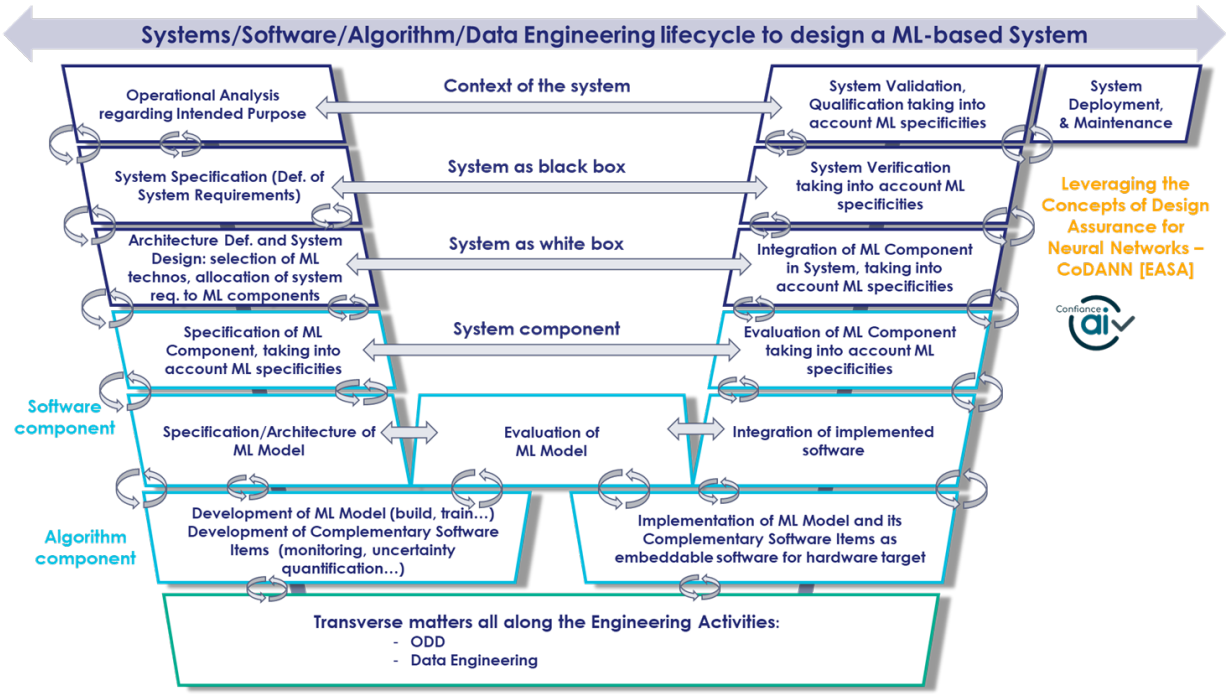

Figure 17: Reference framework for responsible and industrial AI engineering defined in the Body of Knowledge of the Confiance.ai program, integrating the end-to-end methodology.

To meet this challenge, France has launched numerous initiatives, such as the ANITI AI cluster[11] with its DEEL project[12] and the Confiance.ai program,[13] which have defined methodologies and tools for industrial and responsible AI engineering, transforming the classic V-cycle life cycle of a system into a W-cycle[14] and covering the following steps.

Problem specification is captured through analysis of the operational domain (OD) and through the intended purpose, meaning the set of functions, tasks, and objectives for which the AI system will be designed. This step includes definition of performance requirements, usage constraints, application limits foreseen by the manufacturer or developer, functional and non-functional requirements, and operational coverage, e.g. the description of the conditions under which the capability is designed to function correctly, including environmental conditions and other domain constraints. This affects data collection and knowledge modeling, as well as the clear definition of Concepts of Operations (ConOps) and the ODD,[15] which specifies the operational context in which the AI system can function safely and in accordance with its intended purpose. The ODD therefore describes environmental parameters, usage scenarios, and technical restrictions.

Data and knowledge acquisition guided by the ODD leads to aggregation of data and knowledge into a homogeneous set of sufficient size and quality–understandable, relevant, reliable, balanced, and so forth. To be usable, however, these data and knowledge generally must be cleaned, organized, and sometimes labeled. In some cases, processing is required to make raw information exploitable. This is the task of data engineering, which may be complemented by knowledge engineering.

Support for designing or parameterizing a learning algorithm follows. Even if a statistical or connectionist algorithm can be designed or selected from an algorithm library, once learning is completed the model is refined using the validation dataset. This may involve modifying or eliminating variables and adjusting model-specific parameters–hyperparameters–until an acceptable level of precision is reached. Implementation on the target hardware platform and/or system can affect technical requirements such as latency, memory space, or energy consumption.

After an acceptable set of hyperparameters has been found and model precision optimized, the model is tested and characterized on a dataset, or even evaluated through formal verification. Evaluation may go beyond functional performance such as accuracy and include metrics related to any other expected performance criterion, such as robustness to noise and/or adversarial attacks.

Finally, it must be demonstrated that integration of a machine-learning or AI component preserves the expected trust properties. A framework for “trustworthy AI systems engineering” must therefore be defined in order to develop strategies for system development and IVVQ–Integration, Verification, Validation, and Qualification.

5.3. Conclusions

Artificial intelligence represents a considerable force multiplier for modern land armies. By increasing analysis, decision, and action capabilities while reducing soldiers’ exposure to risk, AI profoundly transforms the conduct of land operations. The rapid evolution of AI technologies also yields promising new applications for land forces.

Integrated multi-domain systems: The future lies in seamless integration of land, air, naval, space, and cyber capabilities within multi-domain combat systems. AI will play a central role in this complex coordination, enabling precise synchronization of military effects across domains.

Collaborative robot swarms: Advances in distributed AI and collaborative learning will enable the deployment of swarms of ground robots capable of operating in a coordinated way for reconnaissance, area-securing, or fire-support missions.

This technological revolution, however, will not occur without addressing challenges.

---

[11] ANITI (2024), Artificial and Natural Intelligence Toulouse Institute.

[12] DEEL (2024), Dependable, Certifiable and Explainable Artificial Intelligence for Critical Systems.

[13] Confiance.ai (2024), the community to accelerate deployment of responsible AI in industrial systems.

[14] EASA (2021), *Concept Paper: First usable guidance for Level 1 machine learning applications*, European Union Aviation Safety Agency.

[15] Intended Purpose refers to the set of functions, tasks, and objectives for which the AI system was designed, including performance requirements, usage constraints, and application limits foreseen by the manufacturer or developer. Operational Design Domain (ODD) defines the specific operating conditions under which an AI system is expected to function safely. Concept of Operations (ConOps), including Operational Domain (OD), describes operating scenarios from the user’s perspective, along with procedures and environment. ConOps, including OD, and ODD are interdependent: ConOps provides operational requirements, while ODD translates those requirements into a multidimensional parameter space that must be covered by validation activities. In summary, the intended purpose sets “what” the system must do, while the ODD defines “where” and under which conditions, and “how” the system can do so while respecting safety and performance requirements.

Strategic advantage will belong to nations able not only to develop these cutting-edge technologies but also to integrate them judiciously into doctrines of employment and organizational structures, while adequately training personnel for their optimal use. In this innovation race, the balance between technological progress and wisdom in the use of these new tools will be the key to success.

France's strategic sovereignty in defense and AI can no longer rely on a passive or fragmented approach. Faced with a deeply changing geopolitical and military context–marked by the emergence of denser, faster, and asymmetric threats; the blurring of boundaries between civil and military; and the rise of state and non-state actors less constrained by ethical or legal norms–France must strengthen its decision-making and operational autonomy. This autonomy, however, does not mean withdrawal or a "do everything ourselves" approach. It requires a balanced strategy combining mastery of critical technologies such as AI, targeted cooperation with trusted partners, and industrial agility to adapt to technological disruption.

The issue is twofold: first, to reduce external dependencies in key domains, whether conventional capabilities–stocks, reversible production, low-cost drones–sovereign artificial intelligence–data, infrastructure, models–or autonomous systems–attribution, resilience, interoperability; and second, to accelerate innovation by accepting the temporary lifting of some constraints that slow experimentation. Today, legal, ethical, and administrative timelines slow operational validation and incremental qualification of solutions, even as adversaries or competitors test, deploy, and improve their systems under real conditions, as shown by the war in Ukraine. To close this gap, it is imperative to create more flexible experimentation frameworks allowing industry and land forces to:

- validate concepts rapidly–3D-printed drones, embedded AI, expendable systems–in controlled but realistic environments;
- iterate on prototypes without waiting for full certification, by adopting a staged approach based on progressive qualification; and
- anticipate dual uses by collaborating with innovative civilian actors while maximizing the sovereignty of sensitive technologies.

The creation of AMIAD and initiatives such as Thales's cortAIx or SafranAI show that France has clearly recognized the problem and some of the solutions.

But moving from analysis to action requires the courage to adopt a systemic approach: federate ecosystems–research, industry, and operational users; simplify processes for critical projects; and accept a calculated level of risk in experimentation, without renouncing essential safeguards in security and ethics. Sovereignty is not decreed; it is built through action, by learning faster than the adversary while preserving the strategic alliances that strengthen collective resilience. The time for pure reflection has passed: what is now required is rapid iteration and field validation.

Dependence on American technologies and standards–software frameworks such as Kubernetes and Docker, cloud infrastructures such as AWS and Microsoft Azure, specialized chips such as NVIDIA GPUs, technical and methodological approaches such as MOSA,[16] or AI platforms such as Dataiku and Palantir–nevertheless represents a major strategic risk for European and French sovereignty. These tools, often subject to extraterritorial regulations such as ITAR or EAR, may at any time become levers of geopolitical pressure: access restrictions, blocking of updates, or worse, exfiltration of sensitive data under the guise of legal compliance.

For France and Europe, the urgency is not only to develop sovereign alternatives–trusted cloud projects, European AI accelerators, or controlled open-source frameworks–but also to rethink the standards themselves. Today, technical standards such as ISO and IEEE and interoperability frameworks such as NATO and EU frameworks are largely influenced, if not dominated, by American and Chinese actors, perpetuating a strategic asymmetry. To break this monopoly, France and its European partners must:

- Invest massively in standards bodies–ISO, IEC, CEN-CENELEC, ETSI–to propose alternative standards, especially in:
    - defense AI, including federated-learning models and algorithm auditability;
    - sovereign cloud infrastructures, enabling interoperability without dependence on AWS or Azure; and
    - critical software architectures, including alternatives to Kubernetes or Docker such as K3s or OpenShift under European control.
- Promote European industrial coalitions to certify ITAR-free solutions, relying on:
    - alliances such as the European Defence Fund and Permanent Structured Cooperation (PESCO); and
    - public-private partnerships to develop resilient supply chains, such as RISC-V chips and open-source AI frameworks including Hugging Face.

[16] Initiated by the U.S. Department of Defense in the 1990s and then adopted by many governments and international organizations, MOSA–Modular Open System Architecture–is an engineering methodology aimed at designing systems composed of interchangeable modules, promoting reuse, interoperability, and evolutionary updating.

In addition, the development of a new discipline dedicated to trustworthy-AI engineering appears to be a strategic necessity for France, in order to meet the technological, economic, and geopolitical challenges of tomorrow, especially for deploying trustworthy AI-based solutions in support of land operations. Although generative and agentic AI currently dominate investment, other subdisciplines–such as symbolic, hybrid, and distributed AI–are relevant for deploying innovative capabilities for land forces and essential for guaranteeing critical properties: reliability, robustness, transparency, and maintainability. These qualities are indispensable not only for establishing trust in AI systems but also for qualifying or approving such systems in a context where technological sovereignty and resilience to cyber threats are becoming major issues. Moreover, as stated above, the French ecosystem currently suffers from a lack of harmonized standards, convergent tools across the AI value chain, and dependence on digital giants. Dedicated engineering would make it possible to federate actors–industry, academia, and start-ups–around common methods; accelerate the development of sovereign and frugal solutions; and strengthen the cybersecurity of systems–an imperative in the face of emerging threats such as deepfakes or attacks on sensitive data. Finally, this new discipline of trustworthy-AI engineering is a true lever for anticipating technological disruptions such as Quantum AI and for positioning France as a European leader in trustworthy AI, in alignment with regulations and initiatives such as the European Trustworthy AI Association[17] and sovereign data spaces.

Finally, talent attractiveness is now a major strategic issue for organizations, particularly in innovative sectors such as defense and critical technologies. Attracting competent, diverse, and motivated profiles not only strengthens team performance and agility but also helps anticipate the challenges mentioned above. This attractiveness, however, is not sufficient on its own: it must be accompanied by continuous investment in AI training and acculturation for all land forces and stakeholders. AI, now expanding rapidly, is also transforming professions, decision-making processes, and collaboration modes. To fully exploit it, its use must be democratized, its tools mastered, and its ethical and operational issues understood. Such an approach guarantees not only team adaptation to technological change but also the sustainability of competitive advantage and organizational resilience amid shifts in the industrial and security landscape. By combining talent attractiveness and collective upskilling, AMIAD and the industrial actors of the DITB position themselves as responsible leaders, able to innovate while guiding their engineers toward excellence and thereby proposing innovative AI-based capabilities for land operations.

6. Annexe : Acronyms

AMIAD	Defense Artificial Intelligence Agency
ATDR	Automatic Target Detection and Recognition
BITD / DITB	Defense Industrial and Technological Base
C2	Command and Control
CNN	Convolutional Neural Networks
CSP	Constraint Programming / Constraint Satisfaction Problem
ConOps	Concept of Operations
DH	Hosting Directorate / Direction des Hébergements
DRI	Detection, Recognition, Identification
EAR	Export Administration Regulations
GINN	Geometric-Informed Neural Network
GNSS	Global Navigation Satellite System

[17] The European Trustworthy AI Association (ETAIA) is a non-profit organization created in 2025 to federate the European ecosystem around responsible and industrial AI. It stems from the Confiance.ai program, funded under France 2030, and aims to facilitate the design, validation, and deployment of reliable, explainable AI systems compliant with European regulations. Its missions are to promote trustworthy AI by providing an engineering methodology and state-of-the-art open-source tools; accelerate compliance with AI Act requirements and other regulatory frameworks; create a European ecosystem bringing together major companies including French industrial firms–Air Liquide, Airbus, KNDS, MBDA, Naval Group, Safran, Sopra Steria, Thales–SMEs, start-ups such as Numalis, research laboratories, universities, regulatory bodies, and certification organizations; and support European autonomy in responsible AI innovation.

GOFAI	Good Old-Fashioned AI
AI	Artificial Intelligence
HMI	Human-Machine Interface
IR	Infrared
ISR	Intelligence, Surveillance, Reconnaissance
ISTAR	Intelligence, Surveillance, Target Acquisition, Reconnaissance
ITAR	International Traffic in Arms Regulations
IVVQ	Integration, Verification, Validation, Qualification
LLM	Large Language Model
M2MC	Multi-Domain, Multi-Field
MDO	Multi-Domain Operation
METOD	Tactical Operational Decision-Making Method
OODA	Observe, Orient, Decide, Act
ORTAC	Operational Resource and Tactical Action Control
PINN	Physics-Informed Neural Network
ROEM	Electromagnetic Intelligence
TRL	Technology Readiness Level

Quand l'excellence devient **VITALE**

Groupement des industries françaises de défense et de sécurité terrestres et aéroterrestres

39 rue Mstislav Rostropovitch
75017 Paris
+33 (0)1 44 14 58 20
contact@gicat.fr

gicat.com