



Artificial Intelligence

For Command and
Control Centers

THALES
Building a future we can all trust

Artificial Intelligence

For Command and Control Centers

With the contribution of: Patricia Besson, Sébastien Chenebeau, Philippe Gasc, Katarzyna Kapusta, Jérôme Kodjabachian, Anthony Larue, Christophe Lebreuche, Aldo Maalouf, Hélion du Mas, Juliette Mattioli, Christophe Meyer, Jaime de Oliveira, Arnaud Samama, Claire Trefoux, Matthieu Walle, Cyrille Guetin

Summary

AI is transforming C2 systems into faster, smarter, and more adaptive tools, essential to cope with the growing complexity of modern conflicts. It enhances decision-making, automates tasks, anticipates threats, and optimizes maneuvers. New generations of C2 integrate AI modules and MLOps chains within tailored architectures. Thales offers robust solutions to address the many challenges posed by AI integration in C2. AI remains an assistant: it supports command without replacing it, strengthening operational efficiency while respecting military hierarchy.

1. What are C2s ?

NATO, in its reference documents such as Allied Joint Publication AJP-3, defines Command and Control (C2) as follows:

"C2 is the exercise of authority and direction by a properly designated individual over assigned and attached forces in the accomplishment of the mission."

C2 encompasses key functions:

- Decision-making: planning, organizing, directing, and controlling military operations.
- Information transfer: orders, reports, and intelligence exchanges.
- C2 systems: infrastructures, processes, networks, and technical means ensuring command efficiency and force coordination.
- Command is a responsibility entrusted to a military officer, associated with the authority necessary to analyze, understand, and assess a situation; to conceive, organize, plan; to issue orders, direct execution, and evaluate the results.
- Control is included within command; it is the authority that directs execution and evaluates the results.

A C2 Information System (C2 IS) supports a business process such as in the French Army or COPD in NATO. It is based on the determination of an effect to achieve and enables the structuring of actions. A C2 IS is connected to sensors (cameras, radars, weather stations, etc.) and weapon systems. It is hosted on various types of infrastructure, ranging from data centers to embedded computers (armored vehicles, aircraft, ships, UxVs). The C2 IS has connectivity adapted to its exchange requirements (range, bandwidth, latency, resilience to kinetic threats, jamming).

The functions of a C2 Information System are numerous:

- Multi-domain intelligence (Land, Air, Space, Sea, Cyber) and multi-spectral (Physical, EM, EO, IR) intelligence.
- Development of the tactical situation

- Understanding and anticipation of enemy maneuvers
- Development of friendly maneuvers
- Order planning, production
- Execution of battle plans, evaluation, correction

In the remainder of the document, the term C2 will refer to the C2 Information System (Command and Control Systems).

The C2 covers the entire battle cycle (plan-prepare-conduct-evaluate), but remains distinct from the weapons system. Depending on the level at which it operates, the C2 must or must not meet real-time constraints (guarantee of execution time, in value and/or standard deviation).

2. How is warfare evolving ?

Since 2010 (notably in Afghanistan), and more specifically since the conflict in Ukraine, the way of conducting war has evolved with the following characteristics:

- Complexity, reflected in the new approach or its NATO equivalent MDO (Multi-Domain Operations)
- Battlefield transparency that lifts the fog of war and brings the real situation closer to the perceived situation
- Acceleration of the tempo, shortening of the OODA loops (Observe, Orient, Decide, Act)
- Robotization (including drones) and the proliferation of actors in the 3D battlespace

In this high-intensity combat, it is crucial to quickly and intelligently arbitrate between the use of mass, low-cost effectors and high-tech effectors, in order to reserve the latter for neutralizing strategic targets or protecting high-value units.

In the face of these new battlefield trends, the soldier's cognitive load increases significantly. Empowering C2 systems with artificial intelligence is necessary to enable units to maintain superiority over the enemy.

3. What is AI for C2 ?

Traditionally, model-based AI and data-based AI are distinguished, although these two approaches tend to converge to form hybrid AI.

Model-based AI is a computerization of human knowledge. Typically, rule-based systems (such as expert systems) fall into this category. Human expertise is modeled. Model-based AI can thus be seen as a somewhat specialized form of algorithmic. It is particularly suited to problems that are difficult to model deterministically.

Data-based AI assumes that knowledge is intrinsically contained in the data.

Machine learning consists, for model-based AI, in refining or completing knowledge by adding new information, and for data-based AI, in better classifying or associating data by annotation, most often by human experts (supervised learning), or by automatic analysis (unsupervised learning). For generative AI, the term self-supervised learning is used because these systems learn by aiming to "guess" elements that they themselves extracted from complete sequences (for example, in text, they try to "guess" the word completing a sentence from which they have extracted one or more words and therefore know by construction).

Reinforcement learning, for both forms of AI—model and data-based—consists in improving performance through experience (trial and error), possibly supervised, with an evaluation function that allows comparing the relative performance of two potential solutions. Genetic programming is a particular form of learning in which populations of potential solutions to a problem are manipulated following the principles of Darwinian evolution (crossover, mutation, selection...).

The hybridization of different forms of AI (mostly learning-based) appears to be the most effective. This hybrid AI allows benefiting from the relative advantages of both approaches (performance,

interpretability, etc.). More generally, the hybrid approach also combines other types of knowledge (mathematical, physical, and so on).

Compared to purely human analysis, AI enables processing immense volumes of data in real time, far exceeding human capacity. Even if a large number of human analysts were "aligned" to perform a similar task, the necessary exchanges between them to converge on a common analysis would render the exercise inefficient.

Learning-based AI can potentially improve its performance over time, unlike traditional algorithms. Properly designed and deployed AI can also detect subtle patterns in data that classical methods miss.

The data processed can be of various types: text, images, videos, sounds, radar, and sonar. AI can interpret complex and fuzzy data that classical statistical approaches struggle to model.

4. What are the benefits of AI in C2 ?

The benefits for operational users of C2 systems are multiple:


- Speed and quality of decision support, enabling gaining the upper hand over the adversary:
 - Massive data processing for better understanding of the situation and threats
 - Proposal of action scenarios
 - Orchestration of the maneuver
- Automation of ancillary and often manual processes, allowing to free up time and reduce soldiers' cognitive load.
- New capabilities, for example, countering FPV (First-Person View) drones made possible by sufficiently shortening the OODA loop (sensor – C2 – effector).
- Cybersecurity, through intrusion detection or identification of malicious behavior.

- 3) Reasoning, inference of intentions, and estimation of areas of interest. This step allows the inference of probable enemy intentions. It aims to answer the questions raised during COPD.

Relying on advanced techniques (hierarchical agglomerative and graph-based clustering, adaptive Kalman filtering, and Bayesian reasoning), the solution transforms heterogeneous raw data into a clear synthesis of the tactical situation. Thus, it facilitates the understanding of operational dynamics, anticipates adversary movements, and optimizes the use of

intelligence assets (notably drone swarms), ensuring a strategic advantage in semi-transparent and highly dynamic battlefields.

To support an upstream process such as the NATO CCIR (Commander's Critical Information Requirements) process, Thales AI enables the handling of large volumes of heterogeneous data over an extended period and thus facilitates understanding of the strategic threat and its evolution (see Caption 3). It detects MDO activity patterns and therefore provides threat anticipation capabilities to better prepare for them.



Caption 3 – Detection of activity patterns to anticipate threats

5.2 Proposal of Friend / ENY Maneuver

As soon as situational awareness (SA) is established and enemy's intentions identified, Thales AI proposes a response which can take the form of a FragO (Fragmentary Order) detailing immediate tactical adjustments, or an Overlay Order that graphically overlays action plans onto an operational map for clear visualization by the units.

Thales AI aggregates diverse data sources (DTM – Digital Terrain Model, bitmap, roads, railways), and generates (Caption 4):

- A 2D "mobility" map per unit. Each pixel has a color indicating mobility levels from easier to more

difficult. This map is built based on a lookup table: for Cavalry, for example, certain road widths are required.

- Combat areas conducive to confrontation: the forest is favorable to infantry but not to cavalry.



Caption 4 – Mobility Map

5.3 Planning, decision support, order drafting

At the heart of decision support lies the challenge of choosing a solution among several, based on a set of conflicting criteria in the sense that no single solution is the best for all criteria. For example, one may consider the allocation of drones to identify detected threats. Sending many drones can reduce response times, but at the expense of the capacity to handle a future priority threat that may be detected later. Once the decision criteria are identified, the most acceptable compromises between these criteria must be specified for the decision-maker. Should we keep 2 drones for future threats or reduce the mission duration by 30 minutes?

Thales MYRIAD tool helps build a multicriteria preference model, faithful to expertise, capable of representing interactions between criteria and explaining decisions in the form of influence scores of criteria on the overall evaluation. This tool was used during the REPMUS 24 exercise for drone planning.

5.4 Technical services

To ensure the quality of services offered to its users, C2 systems rely on sometimes complex technical services, some of which can benefit from AI. For example:

- Creation and maintenance of interoperability gateways to translate messages communicated within the framework of multi-domain or federated operations—that is, between actors using different data models. Without fully automating the process, AI assistants for specification processing and software coding help reduce the update time of gateways while requiring a lower technical skill level from operators.
- Adjustment of data compression levels for transmission according to the constraints of computing, memory, and network resources. When possible, deploying AI on the sender side allows for more or less aggressive compression of complex data to reduce bandwidth consumption before reconstruction on the receiver side, with fidelity that can be chosen based on the criticality of the information.



Caption 5 – Semantic segmentation and compression

Above Caption 5 presents a result obtained on a mobile platform (such as an NVIDIA Jetson).

- Top left: the original image
- Top right: a semantic segmentation before transmission

At the bottom: the semantic masks are strongly compressed while preserving the pixels of regions of interest.

- Bottom left: cars are preserved

- Bottom right: cars, people, and buildings are lightly compressed, the rest is heavily compressed.

Between the two images at the top, compression is by a factor of 100.

- Exploitation of technical logs from a set of heterogeneous subsystems to search for root causes. The use of AI assistants simplifies and accelerates data mining by interpreting queries made directly in natural language, utilizing knowledge and procedures derived from available technical manuals, and correlating information from multiple different sources, in order to save technicians time on long and tedious tasks and allow them to focus on high value-added analysis.

6. Architectural Principles of AI-Powered Motorized C2

An AI-powered motorized C2 is architected differently and develops functions that do not exist in a C2 without AI.


- The exploitation of large heterogeneous data requires an architecture that optimizes access performance to storage and computing resources by eliminating unnecessary interfaces and virtualization layers.
- Data access must be secured with conventional protection mechanisms (rights management, encryption of data flows, etc.), ABAC (Attribute Based Access Control) policies, and DCS (Data Centric Security) standards for data exchange between systems.
- The AI-powered motorized C2 must provide tools based on open standards to enable the coexistence of applications using different technologies, to administer and supervise resources among different users and applications, while ensuring overall resilience and secure data access.
- AI models must be regularly updated, so the C2 must include an MLOps (Machine Learning Operations) workshop, usable either in headquarters or in the field, to enable the preparation of new training data and tests, their annotation, the training of AI models with this new

data, verification of the updated models' performance, and then deployment.

- The human-machine interface must allow querying in natural language (e.g., "alert me if enemy units maneuver"), triggering an AI agent that scans available information sources, analyzes the content, and sends real-time alerts to the user. The C2 must therefore be architected with a technical framework equipped with AI agents dedicated to the various C2 functions.
- Finally, the AI models must explain the sources used to produce the decision, describe the decision tree, and present a reliable confidence score for the result.

Consequently, the architecture of the AI-powered motorized C2 is as follows:

- The OS directly deployed on the infrastructure, with the container and VM orchestrator ensuring both system scalability and resilience.
- Technical services for storage, security, and data access control (RBAC, ABAC, DCS).
- Services dedicated to AI and data management: ingestion and preparation of heterogeneous data in both streaming and batch modes, data lineage (a process that allows visualization and tracking of the data lifecycle, detailing transformations, processing, and reuse), Data Lab for data exploration, Data Viz for data visualization, and an MLOps workshop.



Caption 6 – High-level architecture of AI-Powered C2

7. Challenges of AI in C2 systems and associated solutions

This chapter describes the various challenges posed by the introduction of AI in C2 systems, and the solutions that Thales provides.

7.1 Lack of data

In the military context, access to large volumes of data is often difficult due to data confidentiality, the limited number of certain types of equipment, or their short lifespan (e.g., munitions). While expert systems and decision trees require relatively little data, deep learning and foundation models are highly data-intensive.

The response Thales provides to this challenge of insufficient training data is multifaceted:

The response that Thales provides to the challenge of lacking training data is multifaceted:

- The use of smaller, less expressive models (the expressivity of an AI represents the range of functions the model can represent), pre-training datasets, or parts of already trained models.
- Physical simulations, generation of simulated data, or hybrid data through diffusion, which allow the production of large training datasets.
- Foundation models on specific sensors are under development. This is the case for satellite images.
- Structured latent space models based on distance or hybrid models with non-AI components or physical components are promising.

7.2 Vulnerability

The aim is to protect against malicious AI deception or, conversely, to attempt to deceive AI-enabled systems (ENY). Deceiving AI can be done in two ways:

- Exploiting AI biases by searching for and using weaknesses in the AI module itself or in the processes built around this AI.
- Introducing biases into the AI upstream, before its deployment.

Let's illustrate the first way of deceiving AI with an example in a Command and Control (C2) targeting process. The ENY conducts targeting on high-value Enemy Military Units (friend) deep in the area, based on intelligence (ELINT, IMINT, OSINT). To do this, it uses a sensor data fusion tool and AI decision-support to accelerate the fire control loop. Its priority targets are divisional command posts and counter-battery radars.

Long-range artillery, ammunition and fuel depots.

The intended effect for the Enemy Military Units (Friend) is to protect their high-value units and to expend the ENY's resources on false targets by deceiving the ENY using the following means:

- Diversion of the ENY AI responsible for data fusion and decision support
- Digital decoys through injection of biased intelligence data into adversary information streams (social networks, ELINT/IMINT databases, recovered documents, or intercepted signals)
- Creation of fake units such as realistic inflatable structures, thermal simulators, electromagnetic emitters, and passive robotics.

To improve the robustness of AI systems in C2 against the aforementioned threats and methods, Thales has established a cortAix team of "Friendly Hackers," who intervene during the design phase of a system and afterwards to analyze existing solutions.

The BattleBox library developed by this team, usable during the C2 engineering phase, gathers AI-specific attack and defense tools, such as an image and voice deepfake detector and a backdoor detector within a model. To test a model before deployment, it offers a collection of attacks that allow rapid robustness testing of the model against adversarial examples and queries aimed at extracting sensitive information. The Friendly Hackers can thus conduct sophisticated penetration testing missions.

In the previous example, the Friendly Hackers' Battle Box could be used for the first two points.

7.3 MLOps

The motorized C2 with AI is changing the MRO (Maintenance in Operational Condition) process. AI models need to be updated when the environment,


equipment, or battlefield behaviors evolve. Ukraine teaches us that the speed of changes on the theater is incompatible with the traditional model of requesting an update, then modifying the models in the factory, and finally delivering a new version. The MLOps workshop addresses this challenge. An MLOps pipeline is a set of integrated tools and technologies designed to simplify and accelerate the sequence of operations that transform data into an ML model.

An MLOps pipeline provides the following benefits:

- Automating costly and complex tasks
- Facilitating development in environments with different security levels
- Delivering and creating ethical, explainable, and reliable AI
- Supporting ML products throughout their entire lifecycle

Four roles are required to execute this sequence of operations:

- The data scientist who prepares, explores, transforms, labels, and versions the data
- The cybersecurity expert who detects data poisoning
- The model scientist who trains the model
- The user expert who evaluates the operational performance of the AI-powered C2 function




Caption 7 – Principles of an MLOps pipeline

Carrying out the MLOps process in the mainland allows for pooling scientific expertise, but the experts are then far from the users. Performing it on-site requires infrastructures such as a Data Hub and specific skills. The best solution will likely depend on the size of the

Artificial Intelligence for command and control centers

operation area, so it is necessary to develop a solution that adapts to operational constraints and deployment:

- A model developed in the mainland can be deployed across different theaters and even onto devices located at the Far Edge.
- A model developed on-site can also be deployed at the Far Edge.
- The model developed on-site will also be transmitted to the mainland to be deployed on other Data Hubs within the theater.



Caption 8 – Adaptability of the solution according to the deployment

7.4 Cyber Security

Like any Information System (IS), an AI-powered C2 IS must undergo an approval decision process, which requires a rigorous approach: defining the IS (functions, usage, architecture, maintenance procedures), risk analysis, definition of security measures, and identification of residual risks.

However, the world of Artificial Intelligence is characterized by very rapid technological changes and the continuous emergence of new components and applications, coupled with ongoing work on production data to maintain algorithm reliability. These requirements, which shape the C2 architecture as outlined in §5, impose specific security constraints. Furthermore, the MoD approval process will necessarily need to adapt to account for the specificity of AI-powered C2 IS, particularly with the introduction of MLOps.

7.5 SWaP (Size, Weight and Power)

The Forces process data as close to the theater as possible, where constraints on size, weight, and power (SWaP) are significant. While expert systems and decision trees have low resource demands, neural networks and large language models (LLMs) are much more demanding.

Given the wide variety of platform SWaP characteristics, solutions are not universal and align with those mentioned in §6.1. Currently, they implement low-demand AI (e.g., expert systems or decision trees) or small AI models on FPGAs or small GPUs. A common practice, notably for drones, is to replace a single large model with a library of smaller, less resource-intensive models, switching between them depending on flight conditions (ground resolution, sensor definition, etc.) or mission types.

Other solutions will rapidly emerge, thanks to research in the civilian sector on autonomous robots, personal assistants, and the development of small language models (SLMs).

7.6 Trust and Qualification

The integration of AI into C2 systems requires a fundamental rethinking of traditional qualification and certification approaches. Trust in these systems can no longer rely solely on conventional verification of deterministic software. It must incorporate the probabilistic and evolving nature of AI, where decisions emerge from continuous learning processes and neural networks whose internal logic remains partially opaque. However, demonstrating that AI can be trusted is essential, otherwise it will be rejected by the user. The trust attributes related to AI particularly concern:

- **Explainability and transparency:** users must be able to understand, at least partially, how the system arrives at its conclusions. AI currently lacks common sense, with the risk of not realizing when its response is absurd.
- **Reproducibility:** AI is based on nonlinear functions and could give a very different answer to a very similar situation. An AI must produce the same conclusion with similar input data.
- **Robustness:** The AI-powered motorized C2 must be able to detect data attacks, as discussed in section 7.2.

- **Monitoring:** a trustworthy AI must be able to estimate the accuracy of its conclusion (confidence score). More generally, AI-specific metrology is implemented to ensure that AI behaves correctly during its operation. Thales develops its own tools, for example the SCIO algorithm (Statistical Confidence from Internal Observations) available on GitHub.

Proving that AI can be trusted is a scientific challenge that is being addressed by the *Confiance.ai* project, whose objective is to provide methods and tools designed to guarantee and demonstrate the proper functioning of AI by adapting system and software engineering practices. The Thales MLOps pipeline, as described in section 7.3, already integrates *Confiance.ai* components.

For more details, the reader can refer to the Thales white paper “The End-to-End Trusted AI Engineering Lifecycle for Critical Systems.”

Thales continues its R&D efforts, particularly in the development of confidence scores. An AI always produces a result along with an associated confidence score. For an AI trained only with images of cats and dogs, if presented with an elephant as input, the SoftMax function at the network’s output would produce a result of “cat” with a confidence score of 0.8 and “dog” with 0.2. The correct confidence score should have been 0 since the AI has never seen an elephant in its training data, but the AI does not know what it does not know. Thales is working on a method that establishes a confidence score related to the distance between the input data and the training dataset, within a framework of rules such as the necessity to establish a confidence score without knowing the training dataset. Depending on the operational use case, Thales’ method derives rules and then combines scoring methods to obtain the most accurate possible confidence score.

7.7 Cognitive load

AI in C2 systems helps accelerate the OODA decision cycle. However, the human in the loop remains essential: the AI proposes, and the soldier disposes. The challenge, therefore, is the cognitive overload of the soldier, overwhelmed by the AI’s analyses and

recommendations delivered at the accelerated pace of warfare.

How can this collaboration between AI and the soldier be best organized and made adaptable?

The use of the STPA (System-Theoretic Process Analysis) method during the C2 engineering phase provides part of the solution. This risk analysis method allows identifying—within processes where information systems, machines, and humans interact—the moments of risk, especially those with lethal potential. The method highlights where and when things could go wrong, then proposes countermeasures, design requirements, and checklists for the actors involved.

This method was employed during the CoHoMa challenge organized by the Army's Battle Lab.

This approach will be all the more effective if it can be fed by a measurement (through a device such as a BCI—Brain Computer Interface) or an estimation (via behavioral analysis AI) of the soldier's state (stress, cognitive load).

7.8 User Adoption

For an AI-driven motorized C2 to be accepted and adopted by operational personnel, several conditions must be met, both technical and human.

- The commander must remain master of the tempo and the decision. The AI proposes, the human disposes. Any recommendation from the AI must be understandable, traceable, and justifiable. Decision trees, confidence scores, and source data must be accessible to the commander.
- The AI C2 must preserve the hierarchical chain. The principle of subsidiarity remains a fundamental pillar of French command. The AI should not bypass the N+1 or N-1 levels in the name of speed, but enhance their effectiveness through controlled and contextualized information dissemination.
- The AI must relieve, not burden. It should free up time and attention for high-value tasks. This implies careful design of interfaces and workflows, adapted to operational constraints.

Several ongoing projects (DGA, ONERA, NATO CWIX, and SCORPION) aim to develop calibrated trust models capable of alerting when an operator disengages excessively or delegates indiscriminately. Some

prototypes even explore real-time estimation of implicit trust based on interactions or physiological data. These still exploratory approaches share a common goal: to keep the human in the loop without overwhelming or blinding them.

After the challenge of too little trust, the next issue will be excessive trust. AI should not seek to persuade but to inform. A machine speaks like a machine and must be perceived as such. Operators do not expect an intelligent agent that interacts, but a specialized, discreet, and efficient assistant that stays in its place. The AI stays behind a button, produces summaries, filters information, reduces mental workload—but does not think tactically, take initiative, or simulate human intuition. AI must not reduce the military commander to a technical operator. It should enhance their capacity to act by valuing their specific skills: strategic discernment, tactical intuition, and moral responsibility. This requires targeted acculturation to command AIs—not to master them technically, but to supervise them lucidly.

Acceptance relies on transparency, simplicity, and the possibility to take back control at any moment. AI should not be a black box or an algorithmic fantasy but an integrated, explainable, robust, and mastered military tool. Successful adoption of an AI-powered C2 depends not only on its technological performance but on its ability to strengthen the military leader's mission by restoring their space to decide, maneuver, and maintain the upper hand amid the increasing complexity of modern engagements.