

cortAix

Artificial Intelligence by THALES

AI for Electromagnetic Warfare

AI for Electromagnetic Warfare

Authors:

Thomas Courtat, Cedric Deveraux, Thierry Deveze, Cyrille Enderli,

Thierry Sfez, François Sausset, Antoine Violleau,

Summary

The acceleration of tactical and strategic decision loops, combined with the increased density of sensors and effectors, makes mastery of the electromagnetic spectrum essential in modern high-intensity conflicts. To achieve this, the integration of artificial intelligence (AI) into electromagnetic warfare (EW), known as “Cognitive EW,” constitutes a major force multiplier.

This document presents the main operational benefits offered by Cognitive EW, its practical applications, and the key challenges related to the design, integration, and lifecycle of AI systems.

1. Introduction

Electromagnetic Warfare is undergoing a major transformation. In just a few decades, it has evolved from a “useful” function to a “priority” function, on par with cybersecurity and information warfare. Lessons learned from recent conflicts, notably in Ukraine, demonstrate that mastering the electromagnetic spectrum is critical for anticipating, detecting, and countering threats.

Faced with increasingly frequent, low-cost adversary threats, “soft-kill” solutions—neutralizing threats through electronic countermeasures rather than physical interception—offer a significant strategic and economic advantage. They help avoid the pitfall of using a million-euro missile to destroy a threat that costs only ten thousand. To maximize this advantage, it is essential to integrate advances in Artificial Intelligence (AI) into EW.

Moreover, the growing instability of the geopolitical context leads to rapid and unpredictable evolutions of threats in an increasingly dense electromagnetic environment. This dynamic requires accelerated adaptation of algorithms to new situations, which exceeds the capabilities of conventional approaches. AI, thanks to its agility and learning capacity, meets this critical need.

Finally, operators face an increasing workload: fewer in number, confronted with complex systems and high training requirements, they must make rapid decisions in tense operational contexts. This reality calls for advanced automation, which AI can provide, while keeping humans in the loop to ensure accountability and trust.

2. The benefits provided

2.1. OPERATIONAL GAINS

The expected benefits of integrating AI into Electromagnetic Warfare lies at two levels: assisting humans and optimizing machine processing chains.

Augmented warfighters

AI aims to increase operators’ ability to manage complex tactical situations involving growing volumes of data. It reduces cognitive load by masking unnecessary complexities, particularly those related to micro-control of electromagnetic sensors and effectors in real time. This assistance must remain under human control, with clearly defined and validated areas of use, to ensure trust and accountability.

AI also provides a major advantage: rapid experience capitalization and dissemination, far more effective than traditional knowledge transfer between crews.

Augmented sensors

AI improves sensor performance without replacing traditional processing chains but by reinforcing them. This hybridization increases overall robustness. It reveals structures in dense and complex electromagnetic environments or in timeframes inaccessible to humans (very fleeting or very long). These capabilities enrich rapid tactical situation building and knowledge bases by leveraging operational experience from various sensor types.

Enhanced collective orchestration

To maximize the effect on the enemy, EW relies on multiplying payload carriers, manned or unmanned. AI can provide autonomy to each vector and orchestrate collaboration among deployed systems. The goal: achieve a global impact on the enemy’s electronic order of battle by exploiting synergies between spatially and temporally distributed payloads, ensuring electromagnetic superiority.

A more frugal use of resources

AI supports the consolidation of recurrent or rare events and the efficient sharing of knowledge, reducing redundancies and operational costs (for example, fewer deployments of expensive platforms). It also enables optimization of the “low-tech / high-tech” mix by leveraging less costly and more energy-efficient platforms, thereby contributing to a positive environmental impact.

2.2. TECHNICAL GAINS

The technical gains brought by AI are often linked to the digitization of all sensors, whether they concern threats or intelligence or protection systems.

High volume data processing

The digitization of EW receivers leads to a drastic increase in the volume of data to be processed. Not only does EW not escape this digitization, but it is inherently particularly impacted. Indeed, it must continuously listen, detect, sort, and track all signals present on the operational theatre.

In the case of radar EW, for example, up to tens of millions of pulses must be taken into account every second, depending on the density of emitters in the field. On the communications EW side, depending on the recorded bandwidth, the volume of data generated on the field rapidly reaches several terabytes, which are difficult for a single operator to exploit.

In these situations, AI is particularly relevant to provide new solutions that maximize the value of the information contained within this massive volume of data.


Automatic adaptation of detection and identification to threat variability

Similar to communications EW systems, modern radar threats are increasingly digitized and easily reprogrammable. This software agility—dynamic radar modulations and modes, adaptive protocols and transmission schemes, frequency hopping—challenges countermeasure approaches that rely solely on a priori knowledge stored in embedded libraries.

Any new threat, or variation of a known but undocumented threat in the libraries, may be poorly or not classified at all, leading to inappropriate response measures.

AI offers here a capability for dynamic classification and discovery, complementing the libraries: novelty detection, learning, clustering, enabling the identification and characterization of out-of-catalogue signals.

By combining the constraints of each domain, ELINT and COMINT, with existing models, AI facilitates on-mission adaptation and the controlled enrichment of libraries, while keeping the operator in the loop to ensure the robustness and traceability of decisions.



Classification problems being the first where AI methods have demonstrated their relevance, let us delve deeper into the operational and technical contributions in the specific case of EW.

First of all, a good waveform classification capability requires using the temporal sequence information of the parameters describing the waveforms, which is not straightforward given the data rate and the necessary statistical summaries.

If the waveforms are not known in advance, examples intercepted during the mission must be used to "learn" them—that is, to describe them according to an automatically established representation mode allowing subsequent classification of new interceptions.

A specificity of radar EW is that these training data are generally incomplete (radar pulses are missed with an unknown probability), intermittent (the duration of each interception is shorter than the duration of the waveform to be learned), and not always free from mixtures with other waveforms—especially in a dense environment.

Not taking these specificities into account when developing new AI models, particularly during training, often leads to inappropriate performance.

On the other hand, ad-hoc but thorough adaptations of existing methods, carefully considering the specifics of EW, allow recognizing unknown waveforms that have been encountered before, simplifying the tactical situation, highlighting information of interest or changes, and ultimately relieving the cognitive load on operators.

Artificial intelligence will also enable significant contributions at the level of technical reference frameworks.

Several types of repositories exist: at the sensor level and at the national level. It is precisely the coexistence of these that makes the approach new. At the sensor level, a technical repository is a library

containing signatures that allow signal recognition. Its main drawback is that it is specific to a given sensor and prevents its signatures from being generic and reusable across other sensors. A national repository exists, based on unique parameters. It provides more of a business perspective, based on the protocol itself. Even if the target is the same between the two aforementioned repositories, its definition will differ. This results in a consistency problem, making it difficult to enrich these valuable databases.

Today, we have tools that enable the translation of targets from the national repository for the benefit of the technical repository. This is a rather complex approach and is not fully comprehensive. AI would make it possible to characterize a signal detected by a sensor and extract its essential technical parameters. These parameters could then be used to enrich the national technical repository and ultimately obtain a complete and coherent database.

To conclude on the point of the variability of threat behaviors, it should be noted that this also appears in their collective tactics.

Thus, integrated air defense systems (IADS) distribute the tasks of surveillance, tracking, and firing, with the objective of revealing shots as late as possible, ideally when their targets are within a No Escape Zone of their Surface-to-Air missile batteries.

AI enables the macro-chronological analysis of emission sequences from the different IADS transmitters and the typology of relative geographic positioning in order to establish survivability corridors where raid passages would be possible. Again, given the importance of the order in which events occur on the battlefield, well-mastered sequential AI models have a particular role to play.

An enhanced tactical situation for a faster and more relevant decision-making process

In intelligence missions, the objective of electromagnetic warfare is not only to confirm the presence of expected elements but also to provide information about new developments or changes in the operational theater.

A typical scenario is that upon returning from a mission, the AI has identified the signals it already knew, but has also isolated and recognized other signals unknown to its learning database. Before the next mission, electromagnetic warfare operators must validate this automatic detection of

new targets or classes by the AI and, if possible, label the unknown signals to integrate them into the learning database.

If this is planned from the model's design stage, retraining on this updated database can be performed on a small subset of the model only, enabling short training cycles (potentially even during the mission), which are automatic and require minimal computing resources.

Moreover, the exploitation of large amounts of data, with progressively increased granularity in signal analysis over recent years, offers new acuity in identifying signals within increasingly dense electromagnetic environments.

After processing the cues generated by the sensors, AI enables analysis of signal fragments within intermittent pulse trains, going as far as electronic signatures unique to each emitter (fingerprinting). This ability to identify more precisely, beyond the introduced dynamic variations, helps build more detailed and robust tactical situations, leading to a more relevant decision-making process.

Enhanced jamming and deception capabilities

AI enables real-time optimization of the jamming signal or of the decoy/deception strategy to be applied to all threats that must be handled simultaneously.

This optimization leverages the available self-protection jamming resources of the platform (for example, frequency bands, power levels), as well as collective protection capabilities (escort-jamming pods, EW payloads positioned ahead of the strike package).

Optimization is also applied to counter the counter-countermeasure techniques deployed by adversary emitters.

3. From design to operators

As we have just detailed, the operational contributions of AI in the field of electromagnetic warfare are potentially game-changing and could transform the way operations are conducted in the future. However, these promises must be made tangible by designing these new processing methods and integrating them into systems and ensuring that their performance is maintained throughout operations.

Given the particularly critical nature of defense applications, it is essential to carefully select, for each application, the appropriate techniques or combination of techniques that guarantee controlled and explainable performance. Figure 1 illustrates the diversity of existing AI approaches.

In the state of the art, one can distinguish the family of data-driven approaches from that based on knowledge. While the first family primarily relies on statistical models optimized with reference data to extract implicit knowledge, the second family is more conducive to system verification, validation, and qualification because it is generally based on explicit representations of human knowledge and automated reasoning mechanisms.

At the intersection of these two approaches are so-called hybrid methods, which aim to improve the robustness, explainability, and maintainability of data-driven approaches by combining them with the knowledge-based AI paradigm.

Focusing on the overall effectiveness of an AI-based solution in terms of sound decision-making, an algorithm engineering pipeline must take into account the operational design domain, data collection and preprocessing, the design and implementation of an algorithm on suitable hardware, and its deployment.

3.1. AI GROUNDED IN DOMAIN AND SENSORS

Data science seeks to define generic concepts and methods relevant to various application domains (image, video, sound). Many generic tools and software are available, and one might be tempted to apply them as-is or almost as-is to a specific business or operational problem.

However, in the field of electromagnetic warfare, the specificity of signals requires deep domain knowledge to design adapted and optimal solutions.

For example, to analyse a wideband signal, an analogy can be made between its spectral representation (presented to operators as a "waterfall") and an image,



Figure 1 - IA family

which allows considering the reuse of AI models optimized for image processing. This approach quickly yields initial results that may appear convincing.

However, when taking into account the signal's specificities — notably the non-equivalence between the time and frequency dimensions (unlike conventional images where both axes are of the same nature), the particular shape factor of communication signal spectra, their density, and sometimes complex interleaving — it is possible to design AI models that are more effective while being simpler and therefore much more computationally efficient.

The illustrations below (Figure 2 and Figure 3) first show a typical wideband spectrum in the HF range, where elementary signals share the time-frequency space and sometimes overlap. To interpret this spectrum, it is necessary to detect and characterize the elementary signals that compose it.

While promising initial results can be obtained using an off-the-shelf object detection algorithm, designing a specialized algorithm leads to more satisfactory outcomes, as demonstrated by the red boxes indicating AI detections.




Figure 3 – Result of the decomposition into elementary signals (red boxes) by an AI

Processing and Deployment Constraints

In electromagnetic warfare, the volumes of data to be processed immediately after signal digitization are often very large (up to ~1 Tbps). Latency constraints (for instance, for self-protection), communication bandwidth limitations, as well as data sensitivity require processing to be performed as close as possible to the sensors. This is even more critical when the sensors are embedded. Therefore, processing




Figure 2 - Illustration of a wideband signal using its spectrogram

must be deployed on constrained targets rather than on powerful remote servers.

This must be taken into account from the design phase of artificial intelligence algorithms so that their complexity matches the computing capabilities of the intended deployment targets. A specific AI optimization is then necessary for final integration. Hence, it is crucial to master the embedded hardware/software pair in order to fully exploit the potential of AI and be able to deploy it at all relevant system levels, particularly as close as possible to the raw signal from the sensors.

3.2. HOW TO INTEGRATE IT

In systems

The first question to ask is: where to appropriately position AI within an existing signal processing chain? Depending on the situation, end-to-end approaches can be relevant when the existing solution is generally suboptimal (for example, in cases too complex to be modelled efficiently) and when the problem can be formalized in a way accessible to a purely AI-based method. This type of approach has proven successful, for instance, in an ELINT context with radar pulse groupings.

In other cases, where only certain parts of the existing processing chain are suboptimal, it is more relevant to limit AI to these submodules, even chaining several AI-based modules as needed. This has, for example, been the most fruitful approach so far in our work on COMINT.

The question of hybridization with existing processing also arises: for robustness reasons, it can be relevant to run the existing processing and AI-based processing in parallel, then merge their results.

Finally, on a more practical level, software and hardware integration sometimes require reconsidering practices, since AI engineering has some specificities:

- Often dedicated hardware accelerators relying on proprietary SDKs,
- For continuous integration, AI non-regression testing that must be functional and requires sensor data, etc.

On the hardware side

Architectures used in electromagnetic warfare traditionally rely on FPGAs immediately downstream of the analog-to-digital conversion for initial low-level

signal processing. Depending on the case, FPGAs can handle higher-level tasks or forward data to more conventional CPUs once the data flow has been reduced after these processing steps.

Integrating AI requires rethinking this paradigm because, although FPGAs can accelerate neural networks, they are not the best candidates in terms of performance per watt and programming simplicity.

CPUs can be suitable for AI algorithms that do not rely on neural networks, but for neural networks specifically, their architecture lacks sufficient parallelism to execute them efficiently.

There are now GPUs particularly well-suited to accelerating neural networks due to their massively parallel architecture. They also enable acceleration of other types of processing, such as signal processing, and their high memory bandwidth allows them to ingest data streams on the order of magnitude mentioned earlier. This paves the way for unified software and hardware architectures, thereby simplifying and facilitating the integration of AI in electromagnetic warfare

More specifically for AI, NPUs have been developed for several years, providing dedicated acceleration for neural networks, with power consumption ranging from a few milliwatts to several hundred watts. Their advantage lies in their performance-per-watt efficiency, at the cost of specialization and a porting difficulty that sits halfway between that of an FPGA and a GPU.

In such cases, this leads to a more complex overall architecture, structured as an FPGA/NPU/CPU triptych, which in some situations is the only way to achieve the performance required for certain electronic warfare applications.

On the software side

Beyond the programming languages used and the existing architectures—which often need to be at least partially adapted in order to integrate AI—the main challenge lies in AI engineering itself, which differs from “traditional” software in several respects.

To integrate AI into a continuous development and integration pipeline, the following specificities must be taken into account:

- Neural network models contain learned weights, usually stored as files. It is important to monitor the versions of these weights as well as their impact on performance through functional tests.

- These models are also associated with what are known as hyperparameters, which define, for example, details of the training process (learning rate, etc.). Again, it is crucial to monitor them because, even with identical model architectures, they determine the learned weights and therefore the model's performance.
- Training relies on datasets that must be catalogued and whose evolution must be tracked. Like hyperparameters, they influence model performance but, above all, affect the model's ability to generalize. Ideally, this capability should also be monitored using dedicated tests within the AI industrialization process.

For operators

Integrating AI also raises the question of how it interacts with humans, particularly with electromagnetic-warfare operators.

For tasks close to raw signal processing—where humans were not previously involved—it remains entirely appropriate to maintain the same approach when using AI-based methods.

Conversely, for higher-level tasks, AI can provide decision support, enhance operator effectiveness, and, above all, enable new capabilities that were previously difficult to envision—such as detecting unknown targets and allowing operators to integrate them into short loops, thereby increasing overall system agility and delivering an operational advantage.

AI can be integrated at various points of the OODA loop ("Observe / Orient / Decide / Act") to support electromagnetic-warfare operators throughout sustained operations.

In embedded electromagnetic-warfare systems, we distinguish between the onboard segment and the ground segment, which differ in mission and timelines. AI can unlock new possibilities in this context—for example, by updating models on the ground segment between missions, using the data collected and transferred from the onboard segment at the end of a mission.

To maintain trust in AI models—initially trained under specific operating conditions—despite evolving operational contexts (as illustrated in recent conflicts, where the pace of emergence of new threats/targets has significantly accelerated, particularly in electronic warfare), model performance must be monitored

during operations, as must confidence in their predictions.

3.3. AI FOR COMINT ANALYSIS SUPPORT: PRACTICAL THALES EXAMPLE

Com*

The Com* tool makes it possible to identify communication signals across different frequency bands. It consists both of an algorithm whose core is an artificial intelligence component and a software interface offering the operator several modes of interaction.

Identification relies on a paradigm in which the signature of a signal is extracted and compared with a set of known signatures stored in an AI-based technical reference database. A signature may be defined from domain-specific parameters or learned directly by the AI engine.

This approach goes beyond the limitations of classical AI classification methods, which assume that all types of entities to be recognized are known at design time and represented by numerous examples in the dataset used to optimize the system.

Com* is scalable: starting from a few typical samples of a new waveform, the operator can add this waveform to the system's identification capabilities. Unlike conventional approaches, this addition is performed with a limited number of examples, directly in production and on constrained computing resources.

Another limitation of classical methods is that AI tends to provide a result—sometimes erroneous—with an often-misleading level of confidence. With Com*, a signal from an unregistered waveform will be categorized as "unknown." A few similar waveforms can be suggested to the user, who will then decide whether to assign the signal to one of them or not. In either case, the signal is placed in an advanced analysis space where the expert operator can visualize signals, enrich them with contextual information, correct possible errors, group them with similar signals, and ultimately add them to the technical reference database.

Without aiming to replace the operator, the AI engine assists them in their domain study by:

- Automatically isolating signals of interest within wideband signals,
- Providing appropriate visualizations and measurements,
- Suggesting similarities between signals,
- Highlighting portions of the signal that are richer in information.

This approach guarantees, by design, that the system does not regress as the identification reference database grows. The signature comparison logic also results in trustworthy AI, inherently protected against external attacks.

KIA

Another example of AI implementation for COMINT is KIA, an AI toolkit that assists in the analysis of large audio databases. These AI tools are specifically tailored to the needs of COMINT operators and optimized for audio signals from electromagnetic sensors, significantly outperforming commercial off-the-shelf components.

KIA operates at two scales:

- **At the database level:** KIA helps the operator quickly explore large volumes of data by highlighting the most relevant recordings through voice activity detection, then associating metadata with them (recording conditions, language, topics covered via automatic transcription). KIA also groups recordings according to various criteria to cross-reference information across multiple files. This prior indexing allows prioritization of elements for analysis and directs files to the relevant operators.
- **At the individual signal level:** KIA acts as an enhanced audio player, allowing the operator to navigate recordings, detect segments of interest, improve signal quality, isolate voices and background noise, automatically transcribe and separate speakers, and ultimately assist in producing listening reports.

The main operational benefit, supported by feedback, is not only improved understanding of communications but especially a significant reduction in operators' cognitive fatigue. One might have assumed that the primary advantage of noise reduction was intelligibility; however, an experienced operator is not particularly hindered by unclear sound.

The key benefit is providing clearer information, which increases endurance and effectiveness over time.

For Thales, the approach is not to assemble “off-the-shelf” components but to design a processing logic that supports the operator's work, reduces cognitive load by anticipating needs without ever attempting to replace the human. AI is only one part of the Com* engine, orchestrated by effective and robust algorithms. The human-machine interaction is at the heart of value creation.

4. Conditions for success

As we have just seen, the field of electromagnetic warfare has many specific characteristics that cannot be ignored if one wants to benefit from the gains that AI can bring.

Domain expertise is essential to develop relevant AI models, as well as representative real-world data to ensure the performance and robustness of these models.

The integration of AI within electromagnetic warfare solutions can only be achieved through the adoption of new architectures, at the system, software, and hardware levels alike. These changes, sometimes radical, must be anticipated alongside the development of new models; otherwise, the latter will fail to fulfill their promises.

Finally, it is for and by the end user—the electromagnetic warfare operator—that these changes will take place: their feedback will be essential for validating AI-based solutions, particularly to ensure a certain level of trust in them. Above all, the adoption of AI will depend on the value that AI developments bring to these operators in their daily work.

5. About Thales

Thales is directly involved in the challenges and opportunities described in this document through its diverse activities. For several years, artificial intelligence has been integrated throughout its product portfolio—particularly in critical systems and defense solutions deployed or intended to be deployed in many countries around the world. This operational experience has enabled Thales’ research and engineering teams to develop AI components, tools, and methodologies adapted to both sovereign and international contexts.

Some of these innovations were mentioned in previous sections to help the ecosystem better understand the complexities of AI used for electromagnetic warfare and the practical solutions that can be implemented.

Today, Thales brings together over 800 AI experts within its internal dedicated organization, cortAIx, created in 2024. These experts are divided into three main teams:

- **Labs:** focused on fundamental research and model development.
- **Factory:** responsible for tool chains, cybersecurity, and critical infrastructure.
- **Sensors:** integrating AI into embedded and edge systems.

Together, these teams ensure that AI capabilities are not only state-of-the-art but also operationally viable, secure, and compliant with sovereignty requirements. Thales’ cortAIx teams are strategically located across France, the United Kingdom, Canada, Singapore, Germany and the United Arab Emirates, reflecting the company’s global footprint and its commitment to supporting regional sovereignty through local expertise.





4, rue de la Verrerie 92190
Meudon FRANCE

Tél. + 33(0)1 57 77 80 00

www.thalesgroup.com

