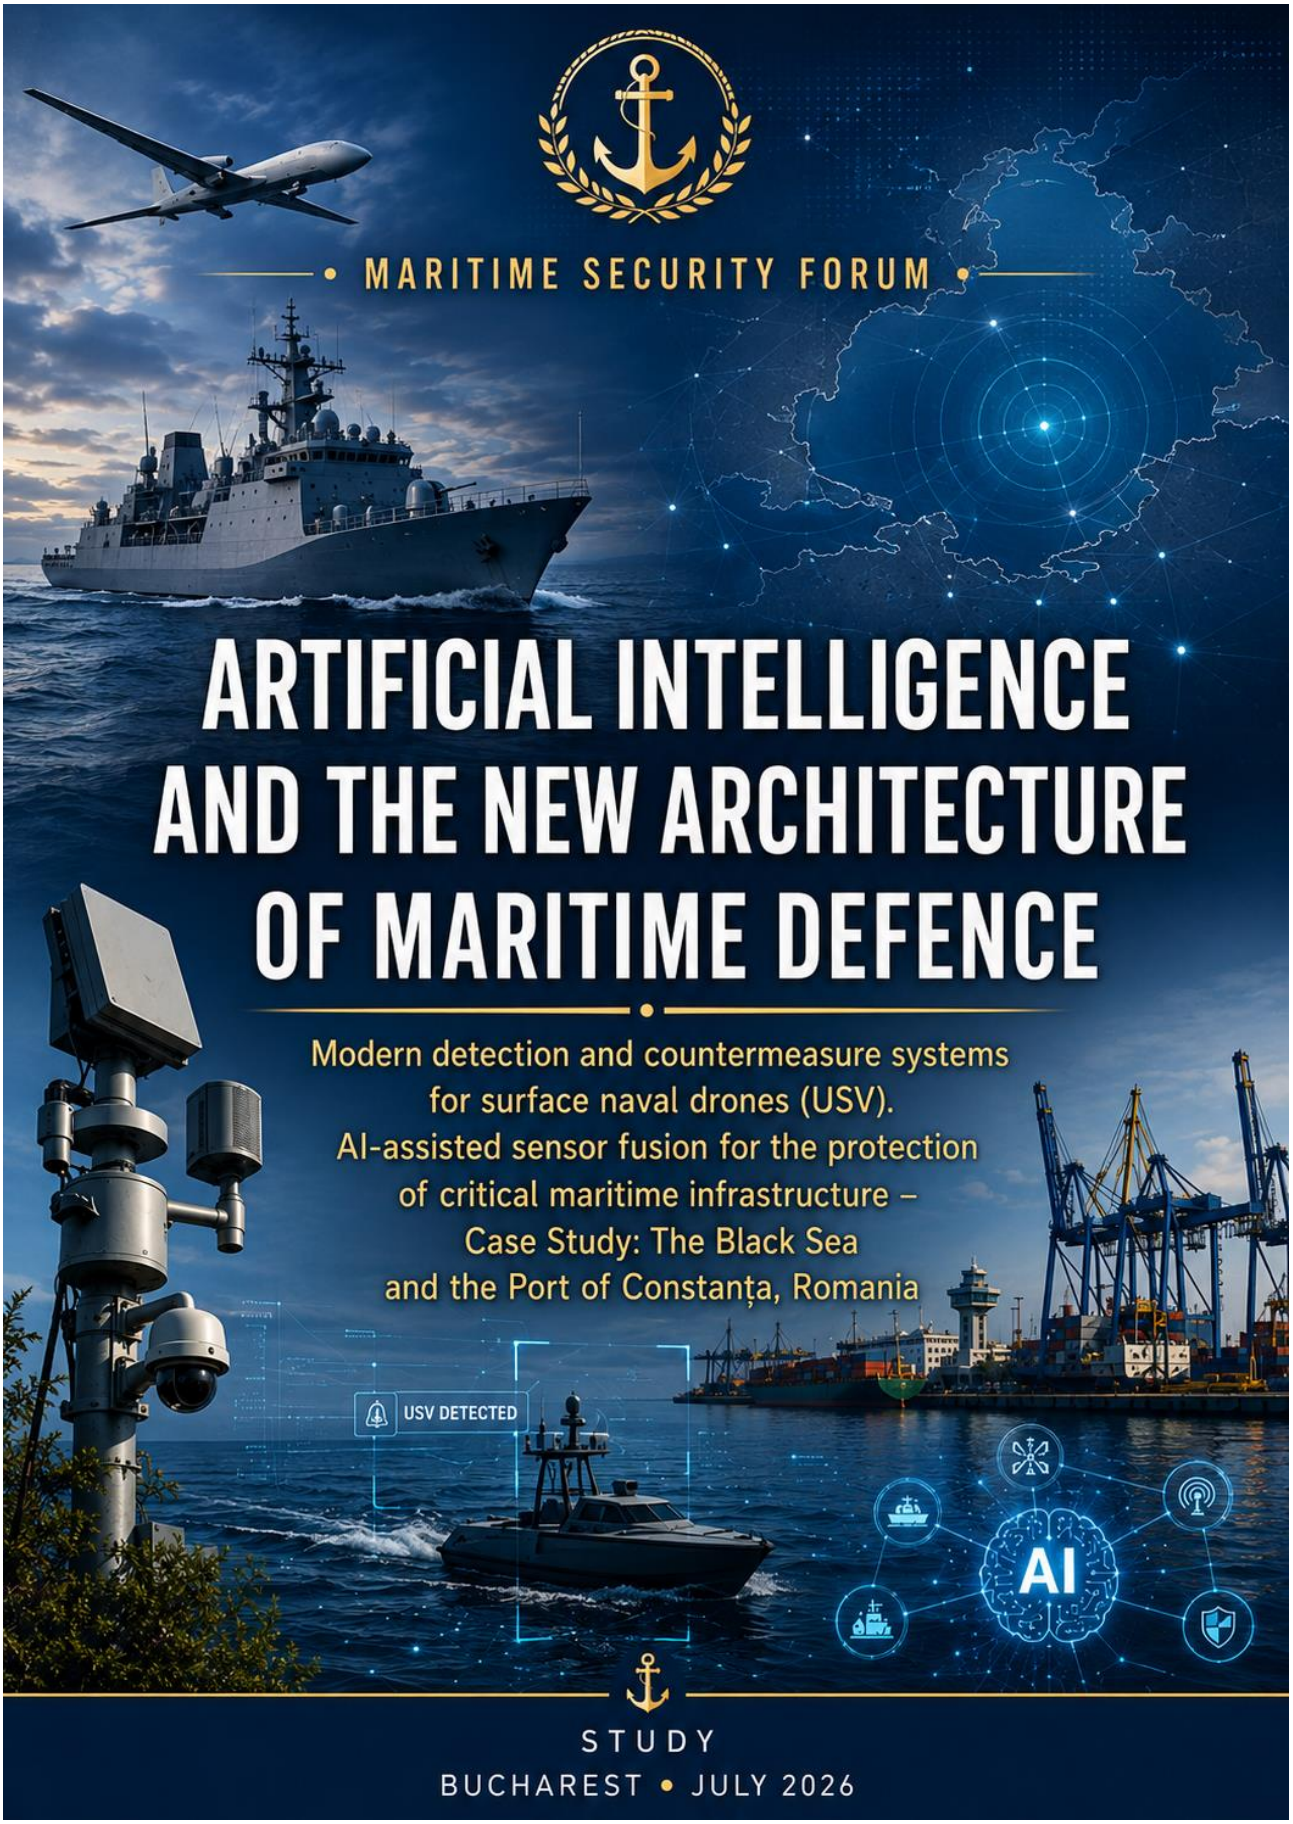
MARITIME SECURITY FORUM
ARTIFICIAL INTELLIGENCE AND THE NEW ARCHITECTURE OF MARITIME DEFENCE
Modern detection and countermeasure systems for surface naval drones (USV). AI-assisted sensor fusion for the protection of critical maritime infrastructure – Case Study: The Black Sea and the Port of Constanța, Romania
USV DETECTED
AI
STUDY
BUCHAREST • JULY 2026

CONTENTS

FOREWORD

Maritime security is undergoing one of the most profound transformations in its contemporary history. The accelerated development of artificial intelligence, the proliferation of autonomous platforms, the expansion of critical maritime infrastructure and the intensification of hybrid threats are altering not only the tools used to protect the maritime domain, but also the very logic of the decision-making process. In this context, strategic advantage can no longer be assessed solely by the number of vessels, the performance of sensors or the capacity for force projection, but by states' ability to rapidly integrate information, technology and human expertise into a coherent system of analysis and decision-making.

This study is based on the conviction that 21st-century maritime security cannot be understood solely through the prism of technological developments, nor exclusively through an analysis of existing legal norms. Current transformations call for an interdisciplinary approach, situated at the intersection of international law, strategic studies, artificial intelligence and the governance of critical infrastructure. From this perspective, this paper aims not only to describe emerging technologies, but also to develop a conceptual framework capable of explaining how these technologies influence the organisation of the decision-making process and the fulfilment of states' obligations in the field of maritime security.

The analysis is based on a review of the specialist literature, official documents produced by international organisations, relevant strategies and doctrines of NATO and the European Union, public reports on the protection of critical maritime infrastructure, as well as information available from open sources (*Open Source Intelligence – OSINT*). The research also drew upon an analysis of recent developments in the Black Sea basin and other strategically significant maritime areas, with a view to identifying the trends shaping the future of maritime security.

One of the study's distinctive features lies in the formulation of a coherent set of original concepts, developed to describe the new cognitive architectures of maritime security. Thus, the paper proposes the concept of **Adaptive Human-Centric Maritime Decision Architecture (AHMDA)**, designed to organise the decision-making process in an environment assisted by artificial intelligence, as well as the concept of **the Adaptive Maritime Security Zone (Adaptive Maritime Security Zone (AMSZ)**), which redefines the maritime security zone as an adaptive space, capable of constantly modifying its configuration based on risk assessments.

To support the decision-making process, **the Dynamic Threat Index**

(DTI), a model for the probabilistic and continuous assessment of maritime threats, whilst **the Maritime Sector AI Readiness Assessment Model** has been developed to assess the degree of artificial intelligence integration

(Maritime AI Readiness Maturity Model (MARMM), a maturity model applicable to modern maritime security architectures. From a legal perspective, the study develops the concept of **Cognitive Due Diligence**, which proposes a reinterpretation of the international duty of care in the context of the use of artificial intelligence, as well as the concept of **Maritime Cognitive Resilience**
(**Maritime Cognitive Resilience (MCR)**), referring to the capacity of a maritime architecture to maintain its essential analytical and decision-making functions in the event of deliberate or accidental degradation of information.

Together, these concepts form what the paper defines as **the vision of maritime cognitive superiority**, according to which strategic advantage is no longer determined exclusively by material superiority, but by the institutional capacity to transform information into operational knowledge, knowledge into legitimate decision-making, and decision-making into effective action. In this vision, artificial intelligence is not viewed as a substitute for the human factor, but as a multiplier of the human capacity for analysis and anticipation, within an architecture grounded in accountability, transparency and respect for international law.

The study does not set out to provide definitive solutions in a field undergoing such rapid evolution. On the contrary, it aims to contribute to the development of an academic and professional debate on the future of maritime security and the role that artificial intelligence may play in a security architecture built around people and the principles of the rule of law.

In a strategic environment characterised by uncertainty, technological competition and rapid transformation, true superiority will not lie with the actor possessing the most autonomous platforms or the most advanced algorithms, but with the one who succeeds in integrating technology, human expertise and legal norms into a system capable of anticipating, deciding and acting legitimately. This is, in essence, the central idea that runs through the entire paper and underpins the concepts and models developed in the following pages.

The concepts developed in this study do not seek to replace the established terminology in the specialist literature on artificial intelligence, maritime security or international law. They represent original conceptual models, constructed by integrating and expanding upon existing doctrinal approaches, with the aim of providing a unified analytical framework for organising the decision-making process and protecting critical maritime infrastructure.

Concept	Does it exist in the specialist literature?	Assessment
Adaptive Human-Centric Maritime Decision Architecture (AHMDA)	No	**Original concept**
Adaptive Maritime Security Zone (AMSZ)	No, in this formulation and definition	**Original concept**
Dynamic Threat Index (DTI)	There are various *threat indices*, but not this model for maritime security	**Original model**
Maritime AI Readiness Maturity Model (MARMM)	No	**Original model**
Cognitive Due Diligence	Not recognised in international law	**Original legal concept**
Maritime Cognitive Resilience (MCR)	There is literature on *cognitive resilience*, but not on this concept as applied to maritime security	**Adaptation and original development**
Maritime Cognitive Superiority	The concept of 'Cognitive Superiority' exists in military doctrine (NATO, USA), but is not defined in this way in the context of maritime security	**Original extension**

ARTIFICIAL INTELLIGENCE AND THE NEW MARITIME DEFENCE ARCHITECTURE

Modern systems for the detection and countermeasures against unmanned surface vessels (USVs). Artificial intelligence-assisted sensor fusion for the protection of critical maritime infrastructure – Case study: the Black Sea and the Port of Constanța

AUTHORS: *Admiral (ret.) PhD Aurel POPA, Rear Admiral (ret.) PhD Sorin LEARSCHI*

Introduction

The transformation of the security environment in recent years demonstrates that the maritime domain has entered a phase profoundly different from that which characterised 20th-century naval warfare. Whilst, in the past, supremacy at sea was determined mainly by the size of fleets, firepower and control of maritime lines of communication, contemporary conflicts highlight the emergence of a new category of threats, characterised by low production costs, high operational flexibility and the ability to generate disproportionate strategic effects.

One of the most significant manifestations of this transformation is the proliferation of autonomous or remotely controlled surface vessels (Unmanned Surface Vehicles – USVs), used for both reconnaissance missions and direct attacks against critical maritime infrastructure. These platforms fundamentally alter the balance between the cost of attack and the cost of defence. Whilst the development and maintenance of a conventional naval system require investments running into hundreds of millions or even billions of dollars, a small naval drone can produce similar operational

effects on a critical target with t a far lower cost. This asymmetry profoundly alters the logic of maritime security and compels states to rethink the way in which surveillance and response systems are designed.

The experience of the conflict in the Black Sea has accelerated this transformation. Since 2022, the repeated use of naval drones against military infrastructure and vessels has demonstrated that such platforms are no longer merely experimental tools, but mature operational components of modern warfare. Commercial ports, naval bases, offshore energy platforms and strategically important maritime routes have become vulnerable to systems capable of operating autonomously or semi-autonomously, navigating with a low radar signature, utilising distributed communications and exploiting the vulnerabilities of conventional surveillance systems.

The Black Sea is today one of the most relevant operational testing grounds for studying these transformations. The geographical characteristics of the basin, the density of energy and commercial infrastructure, the simultaneous presence of the interests of NATO, the European Union and the Russian Federation, as well as the intensive use of electronic warfare, transform this region into an arena where technological developments are tested under real-world conflict conditions. For Romania, these developments are not merely external trends, but realities with a direct impact on national security. The Port of Constanța, the largest port on the Black Sea and one of Europe's main logistics hubs, is home to energy, commercial and military infrastructure, the protection of which is a strategic priority.

The incident that occurred in June 2026 in the Port of Constanța highlighted the fact that traditional maritime surveillance systems, designed to identify conventional vessels and control maritime traffic, face significant difficulties in detecting small naval platforms built from composite materials, with a very low radar signature and capable of navigating partially submerged. This incident demonstrated that the fundamental problem lies not merely in the individual performance of a particular radar or sensor, but in the need for a paradigm shift regarding the entire process of identifying, classifying and neutralising maritime threats.

In this context, artificial intelligence is becoming an essential element of the new maritime security architecture. Its role is not limited to the automation of existing processes, but lies in its ability to integrate vast volumes of data from diverse sources and to generate, within an extremely short timeframe, a coherent operational picture. The fusion of data from maritime radars, electro-optical sensors, thermal cameras, hydrophones, sonars, aerial drones and autonomous patrol vehicles enables a significant reduction in the time required to identify a threat and increases the likelihood of detecting low-signature targets.

The concept of 'sensor fusion' is thus one of the most important areas of development for modern maritime defence systems. Instead of a logic based on each sensor operating independently, the new generation of systems utilises machine learning algorithms and artificial intelligence models capable of simultaneously comparing radar data, visual imagery, thermal signatures, acoustic information and behavioural patterns of targets. The result is the creation of a Common Operational Picture that is far more accurate than that obtained by using each subsystem separately.

At the same time, the expanding use of artificial intelligence presents new challenges. Complex digital architectures themselves become targets for cyber operations, spoofing attacks, the manipulation of sensor data or the compromise of classification algorithms. Consequently, cyber resilience must be viewed as an intrinsic component of modern maritime security and not as a separate domain. Protecting critical infrastructure equally involves safeguarding IT networks, communications and automated decision-making processes.

This study aims to analyse the technological architecture required for the detection and countermeasures against surface naval drones through the use of artificial intelligence and multi-sensor fusion. Drawing on the lessons learnt from the incident in the Port of Constanța and the operational experience gained in the Black Sea, this paper examines the main categories of sensors, command and control software platforms, the integration of autonomous patrol vehicles, associated cyber vulnerabilities and the operational implications for the protection of critical maritime e infrastructure. The analysis aims not only to present existing technological solutions, but also to

highlight how these are fundamentally changing the concept of maritime defence, shifting the focus from post-detection reaction towards anticipatory identification, automatic classification and an integrated response within a cyber-physical ecosystem in which humans remain the final decision-makers.

CHAPTER 1

1.1. The evolution of the threat posed by naval drones in contemporary conflicts

The last decade has witnessed one of the most profound transformations in the maritime domain since the advent of guided anti-ship missiles. Whilst, in the post-war period, naval superiority was associated almost exclusively with fleet tonnage, the number of combat platforms and air-sea dominance, recent conflicts demonstrate that strategic advantage is increasingly being determined by the integration of autonomous systems, artificial intelligence and digital technologies into maritime operations. In this new context, unmanned surface vehicles (USVs) are no longer merely experimental platforms intended for research or surveillance, but systems capable of producing tactical and strategic effects disproportionate to their size and cost.

This development reflects a fundamental shift in the balance between attack and defence. For decades, the protection of maritime infrastructure has been based on the premise that threats would come from conventional naval platforms – military ships, submarines or aircraft – which could be detected using maritime radars, Automatic Identification Systems (AIS), optical observation and aerial surveillance. The emergence of naval drones is radically changing this paradigm. These small platforms, built from composite materials, with a low radar signature and the ability to navigate partially submerged, significantly reduce the effectiveness of sensors designed to detect conventional vessels and are forcing states to rethink their entire surveillance and defence architecture.

Currently, the development of USVs is being driven by the convergence of several technological fields. The miniaturisation of sensors, advances in high-energy-density batteries, the development of satellite communications systems and distributed networks, the integration of artificial intelligence algorithms, and the falling costs of commercial components are enabling the construction of platforms capable of carrying out complex missions without a human presence on board. In many situations, the total cost of a naval drone represents only a fraction of the cost of the munitions required to neutralise it or of the costs incurred in providing permanent protection for a critical target.

This economic asymmetry has significant strategic implications. Rather than engaging in direct confrontation between comparable military platforms, states and non-state actors can use relatively inexpensive systems to threaten infrastructure of exceptional value, such as oil terminals, offshore platforms, submarine communications cables, energy pipelines, port facilities or commercial vessels. Thus, the objective of the attack is no longer necessarily the complete destruction of the infrastructure, but rather the disruption of its operation, the increase in security costs, the disruption of trade flows, and the creation of a disproportionate psychological impact on the economic and political environment.

From a military perspective, naval drones offer an additional advantage due to the difficulty of rapidly attributing responsibility. In many situations, identifying the platform's origin, the operator or the chain of command requires complex investigations, which prolongs the time needed to mount a political or military response. In the maritime domain, where freedom of navigation and the open nature of the environment complicate the identification of a vessel's true intentions, this ambiguity can delay the decision-making process and create operational windows favourable to the aggressor.

The conflict in the Black Sea has demonstrated, in an unprecedented manner, the operational potential of these systems. Since 2022, the repeated use of naval drones in operations targeting military infrastructure and vessels has highlighted the fact that they can carry out reconnaissance,

surveillance, electronic warfare and direct attack missions, sometimes in combination with aerial drones, cyber systems and information operations. Instead of isolated actions, the concept of multi-domain operations is becoming increasingly clear, in which autonomous platforms operate simultaneously in the maritime, air, electromagnetic and cyber domains.

A defining feature of these operations is their integration into a complex digital ecosystem. Modern naval drones do not operate independently, but as elements of a distributed architecture that includes satellites, aerial drones, coastal sensors, communications networks and command centres assisted by artificial intelligence. Consequently, the success of an operation does not depend solely on the technical performance of the maritime platform, but on the entire system's ability to collect, transmit, analyse and utilise information within a very short timeframe. This process dramatically reduces the decision-making cycle and transforms the speed of information processing into a strategic factor comparable to mobility or firepower.

In this new operational environment, port infrastructure takes on unprecedented strategic significance. Modern ports are no longer merely logistical hubs for the handling of goods, but critical nodes within global supply chains, energy security and military mobility. Oil terminals, liquefied natural gas (LNG) facilities, command centres, ammunition depots, electricity grids, communications cables and port management information systems together form an interdependent ecosystem, in which damage to a single component can have a knock-on effect on the entire system.

This reality is particularly relevant to the Black Sea region. Since the outbreak of the war in Ukraine, the strategic importance of the Port of Constanța has increased significantly, with the port becoming one of the main logistical hubs for Ukrainian agricultural exports, allied military mobility and regional energy connectivity. The increase in traffic volume, the diversification of critical infrastructure and the intensification of military activities inevitably heighten the attractiveness of this target for sabotage operations or asymmetric attacks.

At the same time, the evolution of threats demonstrates that the protection of maritime infrastructure can no longer be approached exclusively from a naval perspective. The contemporary maritime domain is characterised by the convergence of the physical and digital dimensions. A naval drone may represent only the visible component of a complex operation that begins with the infiltration of computer networks, continues with the jamming of communications, the manipulation of data streams from sensors and the compromise of the decision-making process, and culminates in the use of the autonomous platform as a kinetic vector. Consequently, effective defence cannot be limited to the physical interception of a vessel, but must include the protection of the entire information chain that underpins the command's operational awareness.

This shift in concept explains why NATO member states and the European Union are accelerating their investment in multi-sensor fusion systems, artificial intelligence, autonomous platforms and distributed command and control networks. The objective is no longer merely to detect a target, but to build a Common Operational Picture capable of integrating information from heterogeneous sources in real time and reducing the time between the emergence of a threat and the adoption of a decision.

From this perspective, artificial intelligence is not merely an auxiliary technology, but the cognitive infrastructure of the new generation of maritime defence systems. Its role is to transform millions of disparate signals into relevant operational information, to identify patterns invisible to human operators, and to support decision-making in an environment characterised by uncertainty, speed and complexity. In the absence of such capabilities, the proliferation of naval drones risks overwhelming the capacity of conventional surveillance systems, and critical maritime infrastructure risks becoming vulnerable to low-cost attacks with major strategic implications.

In this context, the analysis of technologies for detecting and countering naval drones is not merely a discussion of the performance of sensors or software algorithms, but an examination of how the very architecture of maritime security is being reconfigured in the 21st century. The following chapter examines this transformation from a technological perspective, presenting how sensor fusion assisted by artificial intelligence can respond to the new challenges posed by the proliferation of autonomous platforms in the maritime domain.

1.2. A shift in the concept of maritime surveillance: from independent sensors to intelligent defence ecosystems

The technological transformations brought about by the proliferation of autonomous systems necessitate a re-evaluation of the traditional concept of maritime surveillance. For several decades, naval security architectures have been built on the premise that each category of sensor fulfils a clearly defined function: radar detects targets on the water's surface, electro-optical cameras visually confirm their identity, AIS systems provide information on the identity and route of commercial vessels, whilst the human operator integrates all this data to make a decision. This model has responded effectively to conventional threats, characterised by large naval platforms, relatively predictable trajectories and consistent radar signatures.

However, the emergence of autonomous naval drones radically alters this logic. At present, the main difficulty no longer lies in a lack of sensors, but in their inability to correctly interpret an extremely complex operational environment, in which hostile objects are deliberately designed to resemble the natural background noise of the marine environment. Thus, the fundamental problem is no longer the detection of a radar echo, but distinguishing it from reflections caused by waves, floating debris, marine life or other natural phenomena that generate similar signals.

This challenge is exacerbated by technological advances in unmanned naval vehicles. Most modern platforms utilise composite materials, such as carbon fibre or glass fibre, which significantly reduce the effective radar cross-section (RCS). At the same time, the hull geometry is designed to minimise electromagnetic reflection, whilst propulsion systems are optimised to reduce acoustic and thermal signatures. Some platforms operate partially submerged, with only the components strictly necessary for navigation and communications exposed above the water. Under these conditions, the signal received by the radar can become comparable to that produced by a simple wave crest.

This reality highlights one of the fundamental limitations of conventional radar systems. In the maritime environment, the water's surface is one of the most challenging areas for processing electromagnetic signals. Waves, foam, precipitation, and variations in temperature and humidity constantly generate multiple reflections, a phenomenon known in the technical literature as *'sea clutter*'. With a traditional radar, each transmitted pulse generates a very large number of echoes, only a small proportion of which actually originate from objects of interest. The operator must determine whether a particular signal represents a vessel, a buoy, a flock of birds or simply a reflection caused by sea conditions.

The problem becomes even more complex when hostile platforms are specifically designed to exploit these limitations. Modern naval drones frequently travel at high speeds, change direction randomly, minimise their time spent within the radar's field of view, and employ movement patterns adapted to hydro-meteorological conditions. Instead of following standard commercial routes, they exploit the interference zones between natural and artificial reflections, reducing the likelihood of rapid identification.

In traditional surveillance architectures, each sensor operates largely independently. The radar transmits the coordinates of a potential target, the video camera attempts visual validation, and the operator manually compares the available information. This process requires sufficient time for analysis and confirmation. In the case of a naval drone travelling at speeds of over 30–40 knots towards an oil terminal or a port's critical infrastructure, the time available may be reduced to just a few minutes. Any delay caused by the sequential verification of information dramatically reduces the likelihood of an effective interception.

These limitations explain why the current approach is moving towards the concept of **an intelligent maritime defence ecosystem**, in which operational value is no longer determined by the individual performance of a specific radar or camera, but by the entire system's ability to simultaneously integrate and interpret information from multiple sources. In this new architecture,

sensors no longer function as independent elements, but as nodes in a distributed information network, permanently connected via digital command and control platforms.

Multisensory fusion is at the heart of this transformation. The concept involves combining data from radars, electro-optical cameras, infrared sensors, active and passive sonars, hydrophones, AIS systems, aerial drones, autonomous maritime vehicles and satellite sources into a single operational picture. Artificial intelligence serves as the mechanism through which these enormous volumes of information are analysed in real time, identifying correlations that would be impossible to detect using conventional methods.

Unlike traditional systems, in which each sensor produces its own conclusion, architectures based on artificial intelligence aim for the probabilistic convergence of information. For example, a weak radar signal may be insufficient to classify a target. However, if the same object is simultaneously associated with a thermal anomaly detected by an infrared camera, an acoustic signature specific to a waterjet engine, and a kinematic pattern incompatible with the natural movement of floating objects, the probability of a naval drone being present increases exponentially. Instead of analysing each clue separately, the system evaluates them as a whole and calculates the confidence level of each hypothesis.

This approach profoundly alters the role of the human operator. Whereas in traditional systems the operator was responsible for manually identifying targets, in the new architectures the operator becomes the final evaluator of the conclusions generated by the algorithms. Artificial intelligence filters millions of signals, eliminates most false alarms, estimates the probability of a threat and proposes response scenarios; however, the decision on the use of force remains the responsibility of the human operator. This model, known as *the 'human-in-the-loop'* principle, is currently the doctrinal standard adopted by most NATO member states for the use of autonomous systems in the military domain.

The paradigm shift does not concern solely the detection process, but the entire operational cycle of maritime defence. Instead of a sequential response – detection, confirmation, decision-making and intervention – the new systems aim to carry out these processes in parallel. Whilst algorithms confirm the nature of a target, autonomous platforms can already be redirected to the area of interest, physical barriers can be placed on standby, and jamming systems can automatically calculate the optimal parameters for neutralising hostile communications. In this way, the total reaction time is reduced from tens of minutes to a few tens of seconds – a difference that can prove decisive in protecting critical infrastructure.

In Romania's case, this change is all the more important given that the existing maritime surveillance systems were designed for an operational context different from the current one. The SCOMAR architecture represents one of the most advanced regional infrastructures for monitoring maritime traffic and for the surveillance of the European Union's external border; however, the rapid development of autonomous vehicles necessitates the expansion of its capabilities to include a new generation of sensors and algorithms. This is not a matter of replacing existing systems, but of transforming them into an intelligent architecture capable of integrating information from distributed sources and generating an operational picture adapted to new forms of conflict.

This transformation forms the technological foundation of contemporary maritime defence and justifies the development of systems based on artificial intelligence, which are analysed in the following chapter. There, we shall examine in detail the main categories of sensors used in the detection of naval drones and how machine learning algorithms enable us to overcome the limitations inherent in each technology when used individually.

1.3. Artificial intelligence as the cognitive infrastructure of modern maritime defence systems

The transformation of modern maritime surveillance systems is not driven solely by the emergence of more advanced sensors, but by the ability to integrate and interpret unprecedented volumes of information within a timeframe that is incompatible with analysis by humans alone. In

this context, artificial intelligence should not be understood as a mere auxiliary technology intended to automate existing processes, but as the cognitive infrastructure that enables the functioning of the entire contemporary maritime defence architecture.

In today's operational environment, every surveillance platform constantly generates vast volumes of data. A coastal radar produces thousands of echoes with every rotation of the antenna; electro-optical and thermal cameras provide continuous high-resolution video streams; active and passive sonars constantly record changes in the acoustic environment; and autonomous maritime and aerial vehicles transmit real-time parameters regarding their position, speed, direction and the status of their own systems. If this information were to be analysed individually by human operators, the decision-making process would become impossible to carry out within the very short timeframe available to intercept a threat.

The fundamental problem is therefore not a lack of information, but an excess of it. In contemporary military literature, this phenomenon is described as **'information overload**', a situation in which the volume of data exceeds the operator's cognitive capacity to quickly identify relevant information. In the case of an attack carried out using high-speed naval drones, the difference between identifying a target in the first few seconds and identifying it after two or three minutes can mean the difference between neutralising the threat at sea and it striking port infrastructure.

Artificial intelligence intervenes precisely to reduce this discrepancy between the amount of information available and human processing capacity. Modern machine learning algorithms do not replace the human operator, but rather extend their analytical capacity by automatically filtering the millions of signals generated by sensors and by identifying those patterns that indicate the existence of a real threat.

In modern maritime security architectures, artificial intelligence performs several operational functions simultaneously.

The first function is **automatic detection**. Algorithms continuously analyse data streams from sensors and identify objects exhibiting characteristics consistent with a naval drone. Unlike traditional methods based on fixed detection thresholds, deep learning models use millions of past examples to recognise complex configurations of radar, optical or acoustic signals.

The second function is **target classification**. Once an object has been identified, the system must determine whether it is a commercial vessel, a pleasure craft, a buoy, a marine animal, a floating object or a hostile platform. At this stage, the estimated size, speed, acceleration, trajectory, thermal signature, radar response, acoustic profile and kinematic behaviour are analysed simultaneously. Each category of data contributes to calculating the probability that the object under analysis poses a threat.

A third essential function is **behavioural prediction**. AI algorithms do not merely analyse a target's current position, but attempt to anticipate its future movement. By comparing the observed trajectory with millions of previously used behavioural models, the system can estimate whether the tracked object is simply transiting or is attempting to intercept a critical target. In practice, this capability allows defensive measures to be initiated before the drone actually enters the vicinity of the protected infrastructure.

This predictive dimension represents one of the most significant differences from previous generations of surveillance systems. Whilst conventional systems responded exclusively to events that had already occurred, artificial intelligence-assisted architectures seek to anticipate the operational intent of the tracked platform. Consequently, the focus is shifting from simply observing the environment to probabilistically assessing its future development.

Another area in which artificial intelligence is bringing about a major change is the reduction of false alarms. The maritime environment is characterised by extremely high variability. Reflections from waves, flocks of seabirds, schools of fish near the surface, temperature variations or floating objects constantly generate signals that are likely to be misinterpreted as threats. In the absence of effective filtering algorithms, operators are exposed to 'alarm fatigue', where the high frequency of false alarms reduces their ability to respond to real events.

Modern artificial intelligence models use advanced pattern recognition techniques to eliminate the majority of these alarms. The systems analyse not only the instantaneous characteristics of an object, but also its behaviour over time. For example, a naval drone generally maintains its speed and direction in a different way to an object carried by currents or waves. Correlating this information allows for a significant increase in detection accuracy.

From a technical perspective, the operation of these systems is based on combining several categories of algorithms. Convolutional neural networks (CNNs) are used to analyse radar and optical images; recurrent models such as LSTM or GRU enable the analysis of time series and the identification of a target's evolution over time; and modern mechanisms such as Transformers and self-attention facilitate the simultaneous integration of information from multiple sources. In recent years, the development of multimodal models has enabled a shift from the independent analysis of each sensor to the integrated assessment of the entire information ecosystem.

This approach is known in the specialist literature as **Multi-Sensor Data Fusion**, representing one of the most important areas of development for contemporary military systems. Rather than assigning each sensor the responsibility of detecting a threat, the intelligent architecture constructs a common operational picture by correlating all available sources. In this way, the individual limitations of each technology are offset by the strengths of the others.

In the case of port infrastructure, this approach is particularly valuable. Modern ports are extremely busy environments, where commercial vessels, tugs, service boats, pleasure craft, floating equipment and industrial installations all coexist. Without analysis supported by artificial intelligence, it becomes increasingly difficult to distinguish quickly between normal activity and the emergence of a real threat as traffic density increases.

However, the use of artificial intelligence does not remove human responsibility from the decision-making process. On the contrary, as algorithms take on a more significant role in identifying threats, it becomes essential to clearly define the limits of their capabilities. In all military doctrines developed by NATO member states, the use of lethal force remains conditional upon the presence of a human element to validate the conclusions of the automated system. This principle, known as **'human-in-the-loop'**, ensures that artificial intelligence supports the decision-making process without replacing the commander's legal and moral responsibility.

Consequently, the role of artificial intelligence in maritime defence is not to replace the human operator, but to transform vast amounts of disparate data into usable operational knowledge. It acts as the element that connects sensors, autonomous platforms, command systems and communications infrastructure into an ecosystem capable of responding to threats with a speed and precision impossible to achieve through conventional methods. It is upon this cognitive foundation that the entire technological architecture analysed in the following chapter—dedicated to modern AI-assisted detection systems—is built.

CHAPTER 2

Modern AI-assisted detection architecture

2.1. Principles of multisensory architecture

Operational experience gained in recent conflicts demonstrates that no single sensor, regardless of its technical performance, can individually ensure the reliable detection of all maritime threats. The marine environment is one of the most challenging areas for surveillance, as each category of sensor is influenced by different physical factors: radar is affected by reflections generated by waves, optical cameras are limited by fog and darkness, infrared sensors are influenced by thermal variations in the environment, whilst acoustic systems vary in performance depending on temperature, salinity and the structure of the water column.

Under these conditions, modern maritime defence architecture does not aim to perfect a single type of sensor, but rather to simultaneously integrate several independent sources of

information into a system capable of generating a unified operational picture. This concept, known in the specialist literature as **multisensor fusion**, is now the design standard for systems intended to protect critical maritime infrastructure.

Multisensor fusion involves the simultaneous collection of information from maritime radars, electro-optical cameras, infrared sensors, hydrophones, active and passive sonars, aerial drones, autonomous maritime vehicles, AIS systems, satellite imagery and operational databases. Artificial intelligence processes all these data streams in real time, eliminating redundancies, identifying contradictions and calculating the probability of a real threat.

The result is not a simple overlay of images from sensors, but the construction of an **integrated operational picture**, in which each piece of information is assessed according to its level of confidence, the operational context and its relationship with other available data.

This approach enables a dramatic reduction in uncertainty, one of the main challenges in the maritime environment.

2.2. Electro-optical and infrared (EO/IR) systems

Within modern defence architectures, electro-optical and infrared systems are the primary means of visually confirming a threat. Whilst radar indicates the presence of an object, EO/IR cameras answer the fundamental question: **what is that object?**

This distinction is essential from an operational perspective. A radar can detect the presence of an echo without being able to determine with certainty whether it belongs to a naval drone, a civilian vessel, a buoy or a group of highly reflective waves. Visual confirmation is indispensable both for reducing false alarms and for providing the legal basis for a decision on the use of force.

Modern EO/IR systems combine very high-resolution video cameras with thermal cameras capable of operating independently of natural lighting. They utilise variable-focus lenses, gyroscopic stabilisation and dedicated graphics processors capable of analysing complex video streams in real time.

Unlike previous generations, modern cameras do not simply transmit images to the operator. Each frame is analysed locally using **Edge Artificial Intelligence** algorithms, meaning that processing takes place directly at the camera level before the information is transmitted to the command centre.

This approach significantly reduces reaction times and communication requirements. Instead of continuously transmitting high-resolution video streams, the system transmits only those sequences in which the algorithm identifies suspicious objects or relevant changes in the operational environment.

From a technical perspective, the algorithms used are convolutional neural networks (CNNs), trained on millions of maritime images collected under a wide range of weather conditions. These models learn to recognise not only the shape of a naval drone, but also the subtle characteristics of its behaviour: the length of its foam wake, its relative position in relation to the waves, the angle of travel ( ), the temperature distribution across the vessel's hull, or the way in which light reflections differ from those generated by natural objects.

Thus, classification is not based solely on the platform's geometry, but on the simultaneous analysis of a very large number of visual and thermal characteristics.

At night, the importance of infrared sensors increases considerably. Even if the platform is constructed from low-radar-reflective materials and utilises visual camouflage systems, the engine, electronic components and power systems inevitably generate detectable thermal emissions. IR cameras enable the identification of these temperature differences even in conditions of complete darkness, light fog or smoke.

An additional advantage is the ability of modern algorithms to perform **automatic tracking**. Once a target has been identified, the system continues to track it without operator intervention, compensating for the movements of the observation platform, wave-induced oscillations and

sudden changes in direction of travel. In the case of a high-speed naval drone, this function enables visual contact to be maintained at all times, even in difficult weather conditions.

Commercial and military platforms developed in recent years demonstrate the maturity of this technology. Solutions such as **SEA.AI**, originally designed to prevent maritime collisions, use deep learning algorithms to automatically classify objects on the water's surface. Other platforms, such as **Lookout+** developed by Greenroom Robotics, go a step further and estimate the target's distance, direction and speed solely on the basis of video analysis, without requiring additional information from radar.

In the military sector, EO/IR systems are integrated directly into command and control platforms, where images are instantly correlated with information from other sensors. Thus, visual confirmation is no longer the final stage of the detection process, but a permanent component of a continuous validation mechanism.

From an operational perspective, electro-optical and infrared systems fulfil five essential functions:

- confirming the existence of a target detected by radar;
- visual classification of the object;
- identifying the platform's design characteristics;
- assessing the target's operational behaviour;
- providing the visual evidence required to authorise intervention.

Consequently, EO/IR cameras should not be regarded merely as observation equipment, but as cognitive sensors capable of transforming visual information into operational data integrated into the decision-making process.

2.3. Intelligent radar – the evolution from echo detection to behaviour interpretation

Radar remains the primary means of monitoring maritime space, but its operational role has changed fundamentally with the advent of autonomous systems and artificial intelligence. Whilst the first generations of radar were designed solely to identify the presence of a target by determining its range and bearing, contemporary systems are designed to interpret the behaviour of detected objects, to assess the level of threat, and to integrate this assessment into a common operational picture used by modern command and control centres.

This transformation is driven primarily by the changing nature of maritime threats. Conventional radars were developed at a time when the main targets were commercial vessels, large military platforms or aircraft. These targets presented large electromagnetic-reflective surfaces and produced radar echoes that were easily distinguishable from the background noise of the marine environment. Modern naval drones radically alter this situation. Constructed from composite materials, with geometric profiles optimised to reduce electromagnetic reflection and capable of navigating whilst partially submerged, they generate an extremely small radar cross-section, in many cases comparable to that of a simple floating object.

The difficulty in detection is not determined solely by the small size of the platform, but also by the physical characteristics of the marine environment. Unlike airspace, where the electromagnetic background is relatively uniform, the sea surface constantly produces complex reflections caused by waves, foam, precipitation, temperature variations, seabirds and other moving objects. This phenomenon, known as *sea clutter*, represents one of the most significant challenges in maritime radar surveillance. In rough seas, the amplitude of reflections produced by waves can become comparable to, or even exceed, that generated by a small naval drone. Consequently, the challenge for modern radar is no longer the detection of a signal, but the identification of that useful signal amidst a vast amount of interference.

For several decades, this problem has been addressed using classical statistical methods. The most widely used solution is the CFAR (*Constant False Alarm Rate*) algorithm, which is used to automatically set a detection threshold based on the average level of background noise surrounding each radar echo. When the signal amplitude exceeds the calculated threshold, the system classifies

it as a possible target. Although this method has proved effective for detecting conventional platforms, it encounters major difficulties in environments characterised by rapid variations in electromagnetic noise. In adverse weather conditions or in areas with intense reflections from the water's surface, the detection threshold must be constantly recalibrated. If it is set too high, objects with a low radar signature are no longer identified. If it is set too low, the system generates a very large number of false alarms, reducing the efficiency of the entire operational process.

Artificial intelligence fundamentally changes this approach because it moves away from evaluating signal amplitude alone and begins to analyse its internal structure and evolution over time. Instead of simply determining whether or not a radar echo is present, modern algorithms attempt to determine whether that echo exhibits the behavioural characteristics specific to a naval drone. Thus, the analysis is no longer limited to the mere existence of an object, but tracks how it moves, accelerates, changes direction, reacts to obstacles or approaches critical infrastructure. Radar thus becomes a tool for behavioural interpretation rather than merely a presence detector.

One of the most significant innovations in this field is the use of the micro-Doppler phenomenon. Whilst the classical Doppler effect allows the speed of a target to be determined by measuring the variation in the frequency of reflected waves, micro-Doppler analysis identifies the very subtle oscillations generated by moving mechanical components. In the case of naval drones, these variations stem from the rotation of the propellers, the operation of waterjet systems, vibrations of the engine shaft, or oscillations of the hull caused by interaction with the waves. Even when the main body of the platform produces very little radar reflection, these microvibrations generate a characteristic spectral signature that can be recognised by artificial intelligence algorithms. In practice, the system identifies the platform's mechanical signature before it can be clearly observed as an individual radar target.

Processing this information involves transforming the radar signal into a two-dimensional representation known as a Range-Doppler map. This simultaneously describes the position and radial velocity of each target and can be interpreted as a complex image of the operational environment. Convolutional neural networks process these representations in a manner similar to digital image analysis, identifying spatial configurations and relationships that are impossible for a human operator to observe. Even when a drone generates only a few distinct points on a radar map, the algorithms can recognise patterns associated with autonomous platforms by analysing the distribution and evolution of these signals.

A key contribution of artificial intelligence is the introduction of the temporal dimension into the analysis process. Whilst conventional radars treat each antenna rotation almost independently, modern models utilise recurrent neural networks, such as LSTM (*Long Short-Term Memory*) or GRU (*Gated Recurrent Unit*), to track a target's behaviour over several successive scanning cycles. This approach enables the construction of a complete behavioural history, from which accelerations, changes of direction, evasive manoeuvres or the systematic approach to a strategic target can be deduced. The radar thus ceases to provide merely a sequence of positions and begins to describe the dynamics of a threat.

Technological progress is also driven by the development of MIMO (*Multiple Input – Multiple Output*) radars, which simultaneously utilise multiple transmitting and receiving antennas to construct three-dimensional representations of the maritime environment. However, the complexity of these systems generates a very large volume of information, which is impossible to process using conventional methods. To address this problem, modern *self-attention* mechanisms, inspired by the architecture of Transformer models, are utilised. These mechanisms enable algorithms to automatically prioritise those signal components that contribute most to target identification, whilst ignoring redundant or irrelevant information. In this way, the system learns to focus its processing resources on the most important features of the operational environment, significantly increasing the probability of detecting low-signature platforms.

Another major advantage of using artificial intelligence is the improvement in *the signal-to-clutter-and-noise ratio (SCNR*). In traditional systems, this optimisation was achieved through fixed mathematical filters, applied uniformly regardless of operating conditions. Nowadays, neural

networks are capable of learning the statistical characteristics of the marine environment based on sea state, wind speed, traffic density or weather conditions, and of adapting signal processing in real time. Consequently, the radar no longer applies the same filtering strategy in every situation, but constantly adjusts its operating parameters to maximise the probability of detection whilst simultaneously reducing the number of false alarms.

All these developments demonstrate that modern radar can no longer be viewed solely as an electromagnetic sensor. Through the integration of artificial intelligence, it becomes a cognitive system capable of continuously learning from the environment in which it operates, interpreting the behaviour of targets and providing the decision-maker not only with information about the existence of an object, but also with a probabilistic assessment of its operational intent. In modern maritime defence architectures, radar is no longer the first step in the detection process, but one of the main providers of operational knowledge, indispensable for constructing an integrated tactical picture and underpinning decisions regarding the protection of critical maritime infrastructure.

2.4. Multisensory fusion – the foundation of the new generation of maritime defence systems

Whilst intelligent radar represents the first level of the modern detection process, the true technological revolution is driven by the ability of current systems to simultaneously integrate information from entirely different sources. In the specialist literature, this approach is known as *multisensor data fusion* and represents one of the most important areas of development in maritime defence systems over the last two decades. The evolution does not consist merely of adding additional sensors, but of transforming the way in which information is generated, analysed and utilised in the decision-making process.

In traditional architectures, each category of sensor operated almost independently. The radar would detect a potential target, optical cameras would attempt to visually confirm its presence, and operators would manually compare the available information before making a decision. This model was sufficient in an environment characterised by conventional threats and relatively generous timeframes for analysis. However, the emergence of autonomous naval drones radically alters these conditions. A platform travelling at over thirty knots can cover more than a kilometre in a very short space of time, which significantly reduces the time available for confirming and classifying a threat. Under these conditions, sequential analysis of information becomes insufficient.

Multisensory fusion addresses this challenge by replacing sequential analysis with simultaneous analysis. Data from radar, electro-optical and infrared cameras, active and passive sonar, hydrophones, AIS systems, GNSS receivers, autonomous maritime and aerial platforms, and satellite sources are collected within a single processing platform. Artificial intelligence does not treat this information as separate streams, but constructs a common probabilistic model of the operational environment. Instead of confirming a target through a series of checks, the system continuously evaluates all available information and calculates the confidence level associated with each operational hypothesis.

This approach fundamentally changes the nature of the detection process. Rather than answering the question of whether or not a particular sensor is observing an object, the system seeks to determine whether all the available data describe the same physical entity and whether it exhibits the characteristics of a threat. A weak radar reflection may have little significance when analysed in isolation, but its significance increases considerably if it is correlated with a thermal anomaly observed in the same area, with an acoustic signature specific to a waterjet engine, and with the absence of a corresponding AIS signal. In this situation, no single sensor on its own provides certainty of the presence of a naval drone, but the convergence of information from independent sources enables the system to achieve a high level of confidence in the threat classification.

From a mathematical perspective, multisensory fusion involves solving a complex data correlation problem. Each sensor operates in its own coordinate system, has its own measurement errors, its own resolution limits and its own processing delays. The command and control platform must determine whether the radar echo detected at a given moment, the thermal image obtained a few seconds later, and the acoustic signal received by a hydrophone located several hundred metres away describe the same target or different objects. This stage, known as *data association*, is one of the most challenging problems facing modern surveillance systems and is the area in which artificial intelligence is making the most significant advances.

The algorithms currently in use no longer focus solely on the geometric correlation of positions. They also analyse the probable behaviour of each target, its movement history, acceleration patterns, variations in radar and thermal signatures, as well as the operational context in which each observation occurs. Thus, the system can determine with a high degree of probability that multiple observations belong to the same platform, even when individual pieces of information are incomplete or affected by measurement errors.

Probabilistic models and sequential filtering algorithms, such as extended Kalman filters, particle filters and Bayesian methods, play an essential role, enabling the continuous updating of estimates of a target's position and behaviour as new information becomes available. In recent years, these methods have been complemented by deep neural networks and multimodal Transformer models, capable of automatically learning the relationships between information from very different sources, without requiring the explicit definition of all correlation rules. This development represents one of the most significant advances in artificial intelligence applied to the field of maritime defence.

From an operational perspective, the major advantage of multisensory fusion lies in the simultaneous reduction of two categories of errors that affect any surveillance system: false alarms and missed detections. A radar may interpret a large wave as a target, whilst a thermal camera may mistakenly identify a temperature difference caused by solar radiation reflected off the water's surface. When analysed separately, these observations can lead to erroneous conclusions. However, when integrated into a single system, the lack of mutual confirmation automatically reduces the confidence level of the alert and, consequently, lowers the probability of an unjustified response. Similarly, a target with a very low radar signature may go undetected by radar, but may be detected by acoustic sensors and subsequently confirmed by thermal imaging, thereby enabling the identification of a threat that would have remained invisible in a system relying on a single type of sensor.

This approach is particularly important for the protection of critical maritime infrastructure. Modern ports are extremely complex environments, characterised by heavy traffic, continuous industrial activity and a very large number of potential sources of interference. In such an environment, the operational value of a system is not determined solely by the technical performance of each sensor, but by the ability of the entire architecture to transform heterogeneous information into a coherent, real-time tactical picture.

The concept of *the Common Operational Picture* (COP) is the practical expression of this philosophy. The Common Operational Picture is not merely a map displaying the positions of detected targets, but a dynamic representation of the entire maritime domain, in which each object is accompanied by information regarding the confidence level of its classification, its movement history, its probability of future movement, its relationships with other targets, and recommendations generated by the artificial intelligence system. In this way, the operator no longer receives merely raw data, but structured operational knowledge, which enables them to focus their attention on those events that pose the greatest risk to the security of the protected infrastructure.

Looking ahead, multisensory fusion will evolve towards the integration of information from an increasingly diverse range of sources. In addition to conventional sensors, systems will utilise predictive meteorological data, commercial and military satellite imagery, expanded AIS information, *open-source intelligence (OSINT*), as well as outputs generated by artificial intelligence systems capable of anticipating the evolution of the tactical situation. Thus, the

objective will no longer be merely to detect a naval drone, but to construct a predictive representation of the entire maritime environment, in which threats can be identified before they become visible through conventional surveillance means.

CHAPTER 3

3.1 Artificial intelligence-assisted command and control architecture: transforming data into operational superiority

The performance of a modern maritime defence system cannot be assessed solely on the basis of the quality of the sensors used or the accuracy of individual detection algorithms. However advanced radars, electro-optical cameras, sonars or autonomous vehicles may be, their operational value remains limited if the information they produce is not rapidly integrated into a coherent command and control process. In contemporary architectures, the decisive advantage no longer lies with the system that collects the most data, but with the one that succeeds in transforming this data into operational knowledge and, subsequently, into a decision taken before the adversary.

This transformation represents one of the most significant consequences of the introduction of artificial intelligence into the field of maritime defence. Whereas in traditional systems command centres primarily served to centralise information and transmit orders to units on the ground, they now function as digital platforms capable of processing millions of observations in real time, establishing links between seemingly unrelated events, and automatically generating assessments of the likely evolution of the tactical situation. In this sense, the command centre ceases to be a mere information hub and becomes the cognitive core of the entire defence architecture.

Conceptually, this evolution can be understood in relation to the OODA (*Observe – Orient – Decide – Act*) decision-making cycle, formulated by US Colonel John Boyd and used today in most Western military doctrines. The model describes the sequence of stages through which a commander perceives the operational environment, interprets the available information, makes a decision and orders its execution. In contemporary conflicts, success depends not only on the quality of each decision, but also on the speed with which this cycle is completed. The actor who manages to observe more quickly, understand the situation more rapidly and react before the adversary gains a decisive operational advantage.

In the field of maritime security, the proliferation of autonomous drones is dramatically shortening the duration of the OODA cycle. An unmanned naval platform travelling at high speeds can reduce the time available for a response to just a few minutes, and sometimes even to a few tens of seconds in the vicinity of critical infrastructure. Under these conditions, the difference between a conventional architecture and one assisted by artificial intelligence is no longer expressed merely in terms of efficiency, but in the actual ability to prevent an attack from taking place.

Artificial intelligence intervenes at every stage of this cycle. During the observation phase, algorithms simultaneously coordinate data streams from radars, electro-optical sensors, thermal cameras, hydrophones, sonar, aerial drones and autonomous maritime vehicles, eliminating redundancies and rapidly identifying relevant information. In the orientation phase, the system correlates the collected data, compares the current situation with previously learnt behavioural models, and estimates the probability that a particular object poses a threat. In the decision-making phase, the platform generates alternative response scenarios, estimates the time remaining until the protected infrastructure is reached, and calculates the likely consequences of each available option. Finally, in the action phase, commands are transmitted simultaneously to autonomous platforms, jamming systems, physical barriers and response teams, significantly reducing the time between the identification of the threat and the initiation of defensive measures.

However, this acceleration of the decision-making cycle does not mean the human factor is eliminated. On the contrary, as systems become more autonomous in gathering and analysing information, the commander’s role focuses on the legal and operational validation of the proposed measures. In modern architectures, artificial intelligence formulates recommendations, estimates probabilities and prioritises available options; however, the decision on the use of force remains, in

accordance with NATO doctrines and the principles of international law, the exclusive responsibility of the competent human authority. This approach is consistent with the *'human-in-the-loop'* concept, which aims to maintain human control over all decisions likely to result in lethal effects or significant legal consequences.

A defining feature of modern command and control platforms is their ability to function as adaptive systems. Unlike previous generations, where operating rules were set in advance and modified only through periodic software updates, the new platforms use algorithms capable of continuously learning from the operational data collected. Every confirmed alarm, every false alarm, every incident and every exercise contributes to the recalibration of predictive models and to improving the system's performance. In this way, the command and control architecture constantly evolves alongside the security environment, reducing vulnerability to new tactics adopted by adversaries.

At the same time, the integration of a large number of sensors and autonomous platforms presents significant technological and organisational challenges. Modern command centres must manage information flows characterised by very large volumes of data, high update rates and varying degrees of reliability. Furthermore, they must operate under conditions of cyber resilience, so that the compromise of a single sensor or IT component does not affect the entire system's ability to continue the decision-making process. For this reason, contemporary architectures utilise mechanisms such as redundancy, functional segmentation, multiple data validation and the geographical distribution of processing centres, thereby reducing the risk of a *single* point *of failure*.

From an operational perspective, the command centre thus becomes the veritable nervous system of the protected maritime infrastructure. Radar, electro-optical cameras, autonomous vehicles, smart barriers and communications systems can be compared to the sensory and executive organs of a complex organism; however, their effective functioning depends on the existence of a structure capable of integrating all available information and coordinating the response of the entire system. Within this architecture, artificial intelligence acts as a cognitive multiplier, reducing the time required to process information and supporting the commander in making decisions based on a comprehensive, real-time operational picture.

In this context, command and control software platforms can no longer be regarded as mere computer applications designed to display data from sensors. They represent complex digital infrastructures in which artificial intelligence, secure communications, predictive analytics and operational resource management converge. The performance of a modern maritime defence architecture is ultimately determined by the efficiency of these platforms, as they transform raw information into an operational advantage and enable the integrated coordination of all available means for the protection of critical maritime infrastructure.

3.2. Modern command and control platforms assisted by artificial intelligence

The development of modern command and control architectures has led to the emergence of a new generation of software platforms that go beyond the traditional function of aggregating information from sensors. These systems use artificial intelligence algorithms to construct a unified digital representation of the operational environment, within which each piece of information is assessed according to the tactical context, the history of observations and the probability of a hostile event occurring. Consequently, contemporary C2 platforms are no longer merely graphical interfaces used by operators, but genuine decision-support systems capable of significantly reducing the time required to identify and classify a threat.

One of the dominant trends in the development of these platforms is the move away from centralised architectures towards distributed ecosystems, in which every sensor, every autonomous vehicle and every command centre simultaneously contributes to building a common operational picture. Instead of a linear flow of information, characteristic of previous generations, current systems utilise mechanisms for continuous collaboration between all components of the

architecture, so that any change detected by a sensor is reflected almost instantly in the tactical picture available to all authorised users.

A relevant example is the **MARS (Maritime Automated Recognition System)** platform, developed by SeeByte. Originally designed to coordinate autonomous underwater vehicles, it has evolved into a platform capable of integrating information from sonar, radar, electro-optical systems and autonomous surface platforms. The system's distinctive feature lies in its use of machine learning algorithms to automatically identify anomalies in the maritime environment. Rather than tracking individual objects exclusively, the platform establishes a statistical model of normal activity in a given area and flags any significant deviation from this behaviour. This approach is particularly useful in port areas, where the density of traffic makes it difficult to quickly identify a hostile platform using conventional methods.

A different philosophy is adopted by the **Synapse** and **Lattice OS** platforms, developed by Anduril Industries. These systems aim to build an integrated digital representation of the operational environment by simultaneously correlating information from a wide variety of sensors: radars, EO/IR cameras, aerial drones, autonomous maritime vehicles, acoustic sensors and robotic platforms. The innovative aspect lies not only in the integration of these sources, but in the use of artificial intelligence to continuously estimate the probable intent of each target. The platform does not merely indicate an object's position, but also provides a dynamic risk assessment, suggesting to the operator which of the existing contacts require immediate attention and which can be monitored passively.

This change is essential from the perspective of maritime operations management. In a commercial port the size of Constanța, where there may be hundreds of radar contacts and dozens of moving vessels simultaneously, the limiting factor is no longer the system's ability to detect objects, but the operator's ability to quickly identify the real threat. Modern platforms reduce this cognitive load by automatically prioritising events and presenting information in a format tailored to the decision-making process.

In the field of visual surveillance, the development of platforms dedicated to image processing represents an equally significant advancement. **SEA.AI Sentry** systems utilise high-resolution electro-optical and thermal cameras, integrated with artificial intelligence processors capable of analysing video streams locally and automatically identifying low-profile objects on the water's surface. Unlike conventional surveillance cameras, these systems do not continuously transmit images to the command centre, but generate alerts only when the algorithms identify patterns consistent with a potential threat. Edge *computing* reduces both response times and the volume of data that needs to be transmitted via the communications infrastructure, which is essential in situations characterised by electromagnetic interference or bandwidth limitations.

The **Spynel/Cyclope** panoramic systems, developed by HGH Infrared Systems, also fall into this category; they utilise thermal sensors with continuous 360-degree coverage. Unlike conventional cameras, these do not track a single object but simultaneously monitor the entire horizon, being capable of automatically detecting and tracking a very large number of moving targets. Combined with AI-assisted classification algorithms, these platforms offer persistent surveillance capabilities, which are extremely useful for protecting port facilities, offshore platforms and maritime energy infrastructure.

Taken together, these solutions highlight a profound conceptual shift. In the past, each system was designed to fulfil a clearly defined function: the radar detected, the camera provided visual confirmation, and the operator made the decision. In current architectures, this functional separation is gradually disappearing. Sensors generate information that is analysed simultaneously by the C2 platform, and the result is no longer a sequence of independent observations, but a unified assessment of the tactical situation. Artificial intelligence acts as an integrating element between all system components, transforming raw data into operational recommendations and significantly reducing the time between the emergence of a threat and the adoption of defensive measures.

From Romania's perspective, integrating such platforms into the SCOMAR architecture would not entail replacing the existing infrastructure, but rather expanding its capabilities by introducing a higher level of analysis and coordination. Coastal radars, optoelectronic systems, aerial drones and autonomous maritime vehicles would continue to perform their specific functions, but their operational value would increase considerably through integration into a common command and control platform. Such an approach would enable a shift from a system predominantly geared towards maritime border surveillance to an architecture capable of ensuring the active protection of critical infrastructure against complex threats posed by autonomous vehicles, cyber operations and multi-domain attacks.

This development is fully in line with the North Atlantic Alliance's current development priorities, which aim to build interoperable digital architectures capable of integrating sensors, autonomous platforms and command systems into a distributed common operational picture. For the Black Sea littoral states, where the time available to react is often limited to a few minutes, superiority will no longer be determined solely by the individual performance of a radar or a vessel, but by the ability of the entire information ecosystem to rapidly transform disparate observations into coherent and coordinated decisions.

3.3. Resilient communications architecture: mesh networks, Edge AI and operations in electromagnetically contested environments

The transformation of modern maritime defence architectures cannot be understood solely in terms of the evolution of sensors or artificial intelligence algorithms. Equally, operational effectiveness depends on the existence of a communications infrastructure capable of functioning in an environment characterised by jamming, electromagnetic interference, cyber-attacks and the deliberate disruption of command and control links. In contemporary conflicts, neutralising communications is often the adversary's primary objective, as even a high-performance detection system becomes useless if information cannot reach the decision-maker in a timely manner.

Operational experience gained in the Black Sea confirms that electronic warfare is no longer a secondary activity in naval operations, but an essential component thereof. In numerous situations, autonomous platforms, aerial drones and navigation systems have been subjected to jamming, *GPS spoofing*, communications interception or disruption of radio links. These actions aim to degrade the adversary's operational awareness and reduce their ability to coordinate prior to launching a kinetic attack. Consequently, communications resilience has become a prerequisite for the functioning of any modern maritime defence architecture.

The traditional model of military communications was based on the existence of fixed command centres and hierarchical links between them and the platforms in the field. In such an architecture, each sensor transmits data to a central node, which processes the information and relays orders to operational units. Although effective under normal conditions, this model has an obvious structural vulnerability: the compromise of the command centre or the interruption of communications with it can paralyse the entire system.

To eliminate this dependency, contemporary architectures are increasingly utilising *mesh* networks. Unlike traditional communications, where each platform relies on a direct link to a single command centre, in a mesh network each node can communicate simultaneously with several neighbouring nodes and relay information to its destination via alternative routes. Every autonomous vehicle, aerial drone, smart buoy, patrol vessel or coastal station functions both as a network user and as a relay for the other components. In this way, single points of failure are eliminated, and the degradation or destruction of a node does not automatically lead to a breakdown in communications.

The operational advantage of this architecture becomes evident in electronic warfare scenarios. If a particular segment of the network is affected by jamming or if a platform is taken out of action, data packets are automatically rerouted via other available nodes, without the intervention of a human operator. Artificial intelligence continuously optimises these routes based on the quality

of radio links, the level of electromagnetic interference and the mobility of the participating platforms, maintaining network functionality even under severe degradation conditions.

In the case of maritime operations, this flexibility is essential. The marine environment is characterised by continuous variations in radio wave propagation, caused by humidity, temperature, sea state and the configuration of the coastline. Furthermore, autonomous platforms frequently operate at great distances from coastal infrastructure, which necessitates the simultaneous use of multiple communication technologies, including satellite links, tactical radio networks and cellular communications where available. The mesh network enables the integration of all these means into a single infrastructure, capable of automatically selecting the optimal route for each data stream. **A** defining feature of the new architectures is the integration of the concept of **Edge Artificial Intelligence**.

In traditional systems, most of the data collected by sensors was transmitted to the command centre for analysis. This approach involves high bandwidth consumption and a significant reliance on the quality of communications. In contrast, edge processing means that preliminary analysis is carried out directly at the level of the platform collecting the information. The autonomous vehicle, smart camera or sensor-equipped buoy processes images, radar signals or acoustic data locally and transmits only the information deemed relevant to the control centre.

This change has significant implications for operational performance. The volume of data transmitted over the network decreases significantly, which reduces the vulnerability of communications to jamming and enables the efficient use of radio links with limited capacity. At the same time, the time required to identify a threat is reduced, as the analysis no longer depends on the continuous transmission of raw data streams to a remote processing centre. Artificial intelligence thus operates directly at the level of the operational platform, transforming each sensor into an element capable of generating operational knowledge rather than merely raw information.

The concept of Edge AI also has an important dimension from the perspective of resilience. In the event that communications with the command centre are temporarily interrupted, the autonomous platform can continue its mission using locally stored artificial intelligence models. It can detect and track targets, avoid obstacles, adapt its patrol route and record all relevant events until communications are restored. In this way, a temporary loss of connectivity does not automatically lead to a loss of operational capability.

In an electromagnetically contested environment, communications protection also involves the adoption of advanced security mechanisms. Modern networks utilise dynamic key cryptography, continuous node authentication and *frequency-hopping* techniques, whereby the transmission frequency is changed hundreds or even thousands of times per second according to a cryptographic sequence known only to authorised participants. This technique significantly reduces the probability of interception and targeted jamming, as the adversary must simultaneously identify and track a very rapid succession of frequencies.

At the same time, resilient architectures involve the use of navigation systems independent of GNSS satellites. The conflict in Ukraine has demonstrated that GPS signals can be jammed or spoofed over very large areas, affecting both military platforms and civilian navigation. Consequently, modern autonomous vehicles combine satellite navigation with inertial navigation systems (INS), Doppler measurements, optical sensors and *visual navigation* algorithms, thereby reducing reliance on a single positioning source.

Another important trend is the development of architectures capable of operating under degraded conditions (*graceful degradation*). Rather than the loss of a single component causing the entire system to fail, modern platforms are designed to gradually reduce their functionality whilst retaining the ability to carry out essential missions. Thus, in the event of a satellite communications failure, the system can continue to operate via the mesh network; if this is affected, the platforms can autonomously carry out previously planned missions; and in the event of further degradation, they can automatically return to a safe area or adopt pre-programmed self-protection procedures. This principle of controlled degradation is now one of the most important design requirements for autonomous systems used in the military.

From the perspective of protecting Romania's critical maritime infrastructure, the development of a resilient communications architecture is at least as important as modernising sensors or acquiring new autonomous platforms. A high-performance detection system quickly loses its usefulness if information cannot be distributed in a timely manner to command centres and response forces. By contrast, an infrastructure based on mesh networks, distributed processing, Edge AI and advanced cyber-security mechanisms enables a coherent operational picture to be maintained even under conditions of intense electronic warfare operations. In a theatre such as the Black Sea, where contestation of the electromagnetic spectrum is already an operational reality, this capability is an indispensable element of any modern maritime defence architecture.

3.4. Autonomous maritime vehicles – from patrol platforms to intelligent nodes in the defence network

Over the past two decades, autonomous maritime vehicles have evolved from experimental platforms designed for oceanographic research and routine surveillance into one of the most important components of modern maritime security architectures. Whilst the first generations of unmanned systems carried out limited missions, based on pre-programmed routes and relatively simple sensors, the development of artificial intelligence, distributed communication systems and autonomous power technologies has radically transformed the role of these platforms. Today, they are no longer viewed solely as patrol vessels, but as intelligent nodes within a complex network of sensors, capable of collecting, processing and distributing operational information in real time.

This change is particularly important for the defence of critical maritime infrastructure. In traditional architectures, surveillance was concentrated near the coastline, where coastal radars and optoelectronic systems could effectively cover the area of interest. However, such an approach assumed that a threat would only be detected once it had entered the range of land-based sensors. In the case of high-speed naval drones, the interval between detection and impact may be too short to mount an effective response. Autonomous vehicles change this dynamic by shifting the surveillance line tens of nautical miles out to sea, where the threat can be identified at a much earlier stage.

In this regard, the autonomous vehicle should not be viewed as an alternative to coastal radar or patrol vessels, but rather as a mobile extension of the entire surveillance system. It operates continuously in the area of interest, collects information via its own sensors and transmits it to the command centre, helping to build a much more comprehensive operational picture than that obtained solely through sensors fixed on the coastline. Its mobility allows for continuous adaptation to changes in the tactical environment and eliminates one of the main limitations of static infrastructure: the inability to rapidly adjust the position of sensors in response to evolving threats.

An additional advantage is the persistent nature of the surveillance. Conventional vessels are constrained by fuel consumption, the need to rotate crews and the high costs of continuous operation. Modern autonomous vehicles utilise renewable energy sources, efficient battery management systems and energy consumption optimisation algorithms, enabling them to carry out long-duration missions without significant interruptions. This capability transforms maritime surveillance from a periodic activity into a permanent presence in the area of interest.

The integration of artificial intelligence amplifies this transformation. Autonomous platforms do not merely transmit raw data to the command centre; they analyse sensor data locally and generate alerts only when they identify relevant anomalies. This distributed processing reduces the load on communications networks and enables the efficient use of available resources even in conditions of electromagnetic interference. At the same time, each vehicle contributes to the continuous improvement of machine learning models by accumulating data on the behaviour of the marine environment, traffic characteristics and observed threat patterns.

From a doctrinal perspective, autonomous vehicles are also changing the relationship between combat platforms and command infrastructure. In the past, most information was collected by manned platforms and transmitted to command centres for analysis. Nowadays, the process is distributed. Autonomous vehicles carry out a significant part of the analysis directly on board,

whilst the command centre receives information that has already been filtered and contextualised. This model reduces the time required for decision-making and allows human operators to focus on assessing real threats, rather than on interpreting a vast volume of raw data.

These developments explain the growing interest shown by NATO member states in integrating autonomous vehicles into maritime surveillance architectures. The aim is not to replace patrol vessels or existing radar systems, but to extend their capabilities by creating a distributed network of mobile sensors, capable of operating continuously in the vicinity of high-risk areas.

3.5. TRITON – an operational model for the defence of critical maritime infrastructure

It is within this technological and doctrinal context that the **TRITON** autonomous vehicle, developed by the American company Ocean Aero, finds its place. Beyond its technical characteristics, the interest in this platform stems from the fact that it illustrates the direction in which modern maritime surveillance systems are evolving. TRITON is not designed as a mere observation platform, but as an intelligent node capable of actively participating in the process of collecting, analysing and distributing operational information.

One of the features that sets this platform apart from most existing autonomous vehicles is its dual architecture, which allows it to operate both on the water's surface and submerged. This capability gives it exceptional operational flexibility. Under normal conditions, the vehicle can navigate on the surface using solar and wind power, ensuring an operational range of weeks or even months. When the tactical situation requires it, it can switch to submerged mode to reduce the likelihood of detection, to avoid adverse weather conditions, or to carry out covert surveillance missions near sensitive infrastructure.

This transition between the two modes of operation is not merely an engineering feat, but meets a fundamental operational requirement: the platform's survival in a contested environment. In contemporary conflicts, autonomous vehicles may themselves become targets of electronic warfare operations or kinetic attacks. The ability to rapidly alter the operational profile reduces the system's vulnerability and enables it to continue its mission even under high-risk conditions.

Energy autonomy is another strategic feature. The combined use of photovoltaic panels, high-capacity batteries and sail-assisted propulsion reduces dependence on logistical infrastructure and enables persistent surveillance missions to be carried out at significantly lower operating costs than those of a conventional vessel. This cost-effectiveness is particularly relevant for the protection of critical maritime infrastructure, where continuous monitoring requires the permanent deployment of significant resources.

Equally important is the platform's integration into a distributed sensor and communications architecture. TRITON does not operate in isolation, but constantly exchanges information with aerial drones, coastal radars, smart buoys and command centres via resilient communications networks. In this configuration, the vehicle's operational value stems not only from its own sensors, but from its ability to contribute to the creation of a common operational picture, in which each platform complements the information provided by the others.

This approach reflects the general direction of development in modern maritime defence architectures. Operational superiority is no longer determined by the individual performance of a single platform, but by the ability of the entire network to function as a unified, adaptive and resilient system. In this sense, TRITON is more than just an autonomous vehicle: it is a concrete example of how artificial intelligence, distributed communications and energy autonomy can be integrated into an architecture designed to protect critical maritime infrastructure.

CHAPTER 4

Designing an integrated defence architecture against naval drones for the Port of Constanța and the Romanian coastline

The technological developments analysed in the previous chapters demonstrate that the protection of critical maritime infrastructure can no longer be achieved through the isolated modernisation of individual pieces of equipment or the acquisition of additional autonomous platforms. Operational efficiency stems from the integration of all components into a unified architecture, capable of detecting, interpreting and neutralising threats before they reach the protected targets. In this context, the design of a system for the Port of Constanța must not begin with the question 'what equipment is needed?', but with the fundamental question 'how should the entire information and operational architecture be organised so that the available reaction time is maximised and the probability of tactical surprise is minimised?'.

This shift in perspective is essential. Traditionally, port security has been built around the concept of perimeter protection. Radars, CCTV cameras and naval patrols were tasked with monitoring the immediate vicinity of the infrastructure, with a response being triggered the moment a threat entered the range of the surveillance systems. The development of autonomous naval drones means this approach is no longer sufficient. Modern platforms can travel at high speeds, have a low radar signature and can follow routes that are difficult to anticipate, which considerably reduces the time between detection and impact. Under these circumstances, effective defence requires moving the surveillance line as far out to sea as possible and transforming the entire maritime area off the port into a permanent surveillance zone.

The architecture proposed in this study is based on the concept of *layered maritime defence*, currently used in the development of critical infrastructure in NATO member states. The principle is that no single category of sensors or platform can, on its own, provide complete protection for a target. Instead, each layer of the system fulfils a specific function, and its effectiveness stems from the integration of these functions within a single command and control architecture.

The first layer must be positioned at a considerable distance from the protected infrastructure and serves as an early warning system. It consists of autonomous maritime vehicles with a long range, deployed along the main access routes to the Romanian coastline. These platforms patrol continuously in pre-defined areas and use acoustic sensors, electro-optical cameras, im thermal imaging and compact radar systems for the early identification of hostile platforms. Unlike conventional naval patrols, these vehicles do not aim to directly intercept threats, but rather to continuously gather intelligence and transmit it to the entire surveillance network. Their role is to extend the system's intelligence horizon and transform the maritime space into an area of persistent observation.

The second layer consists of fixed coastal surveillance infrastructure. In Romania's case, this function is primarily fulfilled by the SCOMAR system, whose capabilities can be expanded through the integration of high-resolution radars, artificial intelligence-assisted thermal cameras and automated target classification systems. The role of this level is not to carry out initial detection, but to validate and supplement the information provided by autonomous platforms at sea. In this way, every identified contact undergoes multi-sensor confirmation before the system recommends a response.

Between these two levels, it is necessary to develop a distributed network of mobile sensors, comprising smart buoys, vertical take-off drones and small autonomous vehicles. These platforms serve to eliminate any blind spots between coastal and offshore sensors, ensuring the continuity of the operational picture. Smart buoys can continuously monitor the acoustic signatures of the marine environment, whilst aerial drones can provide rapid visual confirmation of contacts identified by the other components of the system. The mobility of these platforms allows for rapid adaptation to changes in the tactical situation and the redeployment of resources according to the likely direction from which a threat may emerge.

At the heart of the entire architecture lies the artificial intelligence-assisted command and control platform. It simultaneously receives information from all levels of the system, correlates it, eliminates redundancies and generates a single operational picture. Rather than each sensor independently transmitting alarms to operators, the platform constructs a probabilistic assessment of each contact, taking into account the movement history, radar characteristics, thermal profile,

acoustic signature and kinematic behaviour of the tracked object. Only after a pre-set confidence threshold is reached does the system recommend the initiation of defensive measures.

One innovative feature we propose is the introduction of the concept **of an** *'Adapt***ive** *Maritime Security Zone*'. Unlike the fixed perimeters currently in use, this zone constantly adjusts its configuration according to the level of risk estimated by artificial intelligence algorithms. During periods of normal activity, surveillance is conducted under a standard regime, focusing on the main access routes to the port. When the system identifies changes in the operational environment – such as an increase in electromagnetic jamming, the appearance of unidentified contacts, changes to trade routes, or alerts from external intelligence sources – the security zone is automatically expanded, and the autonomous platforms are redeployed to reinforce sectors deemed vulnerable. Thus, the architecture no longer reacts solely to the emergence of a threat, but proactively adapts its configuration based on continuous risk assessment.

In this configuration, time becomes an operational resource that can be managed through intelligent system design. If a naval drone is detected twenty nautical miles from the protected infrastructure, the command centre has a significantly longer window to assess the situation, mobilise resources and coordinate the response. This additional time allows for multi-sensor confirmation of the contact, reduces the likelihood of false alarms and enables the selection of the most effective neutralisation measure. Consequently, the primary objective of the architecture is not merely to increase the probability of detection, but to extend the time window available for decision-making.

Such a model is particularly well-suited to the Romanian coastline. The relatively short length of the coastline, the existence of a limited number of strategic ports and the concentration of energy infrastructure near Constanța allow for the development of a distributed architecture without the very high costs associated with similar systems designed for extensive coastlines. Integrating existing infrastructure with autonomous platforms and artificial intelligence systems would enable the transformation of the current surveillance model into a digital ecosystem capable of responding to the specific threats posed by contemporary maritime conflicts.

Looking ahead, this architecture could form the basis for developing a national defence concept against autonomous maritime vehicles, compatible both with NATO's interoperability requirements and with the European Union's objectives regarding the protection of critical infrastructure and the strengthening of maritime resilience. In a context characterised by the proliferation of unmanned platforms and the intensification of electronic warfare, the strategic advantage will no longer lie exclusively with the state possessing the most advanced platforms, but with that which succeeds in integrating information, technology and the human factor into an adaptive architecture, capable of anticipating and managing threats before they affect protected infrastructure.

4.1. Adaptive Maritime Security Zone (AMSZ): a new paradigm for the organisation of maritime security space

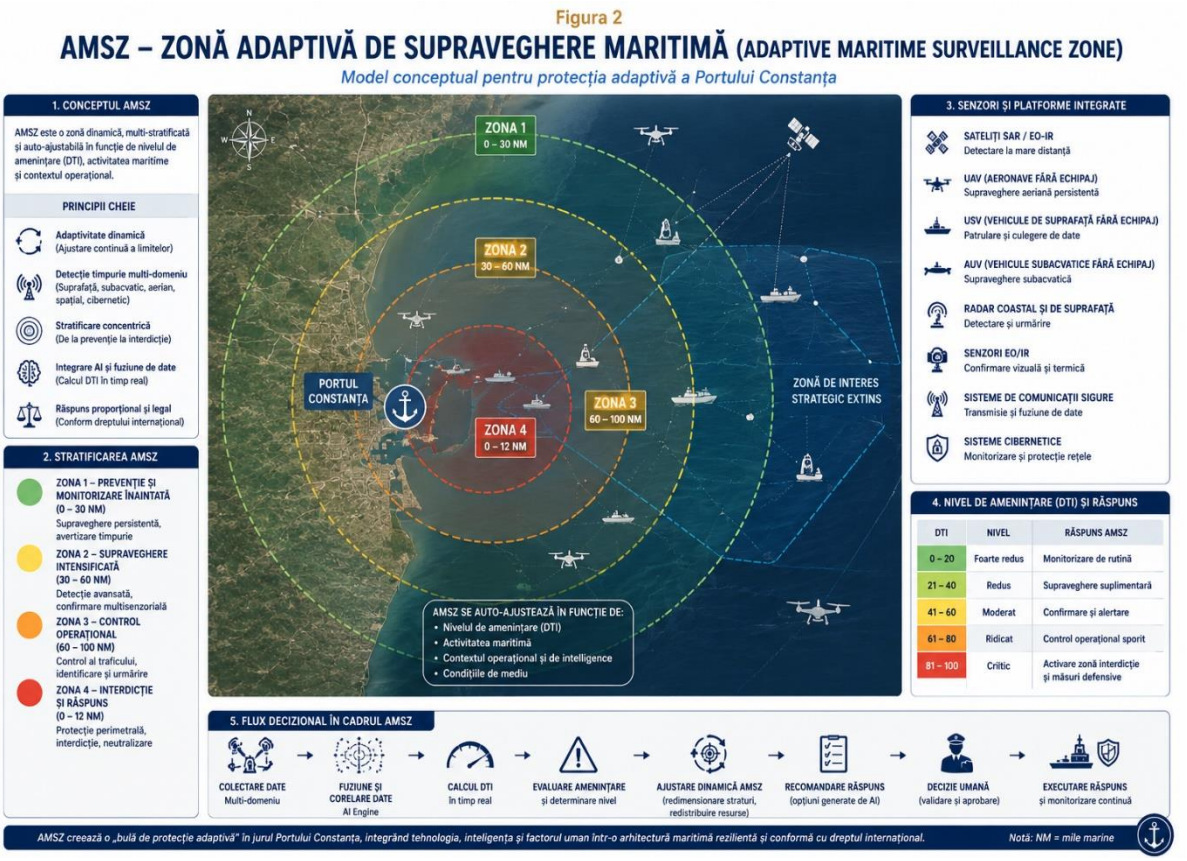

The protection of critical maritime infrastructure has traditionally been based on geographically defined security zones characterised by relatively static defensive measures. Ports, energy terminals, offshore platforms and other essential infrastructure are protected by fixed perimeters, naval patrols, radar surveillance and standardised access control procedures. This approach was appropriate in a context where the main threats were posed by conventional platforms, with relatively predictable trajectories and reaction times long enough to allow for the implementation of defensive measures.

However, the changes brought about by the proliferation of autonomous systems, the use of artificial intelligence and the development of hybrid threats are altering the very foundations of this model. Unmanned naval vessels, aerial drones, coordinated attacks on undersea infrastructure, cyber operations and electronic warfare are creating an operational environment characterised by high mobility, uncertainty and a drastic reduction in the time available for analysis and response. Under these circumstances, the traditional concept of a security zone, defined by fixed geographical boundaries, becomes insufficient for managing emerging threats.

Building on this observation, this study proposes the concept of **the Adaptive Maritime Security Zone (AMSZ),** defined as **a spatio-temporal security architecture in which operational boundaries, sensor deployment, the positioning of autonomous platforms and the level of protective measures are continuously reconfigured based on a predictive risk assessment carried out using artificial intelligence**.

The AMSZ concept is based on the idea that the security space should not be viewed as a static perimeter, but as an operational entity undergoing constant transformation. The size, shape and density of the defensive system are not determined prior to the operation, but result from the continuous analysis of the maritime environment and the anticipation of potential threat behaviour. From this perspective, the security zone becomes an adaptive system, capable of altering its configuration before the threat reaches the protected infrastructure.

The fundamental difference from classical architectures lies in the fact that the AMSZ does not react exclusively to events that have already occurred, but seeks to continuously adjust the defensive posture according to the probability of incidents occurring. Artificial intelligence simultaneously integrates information from radars, AIS systems, electro-optical sensors, sonar, satellite imagery, autonomous platforms and open-source intelligence, constructing a dynamic risk assessment for each sector of the maritime domain. Based on this assessment, the system recommends redeploying sensors, altering the routes of autonomous platforms, stepping up surveillance in certain sectors or activating additional protective measures.

In this architecture, the maritime domain is no longer organised solely according to geographical criteria, but also according to the temporal dimension of risk. Thus, two sectors located at the same distance from the protected infrastructure may be subject to different levels of surveillance if the algorithms estimate different probabilities regarding the evolution of threats. The security zone thus becomes an expression of anticipated risk and not merely of geographical position.

A key element of the AMSZ concept is its multi-level nature. Instead of a single protective perimeter, the architecture operates by overlapping several operational levels that adapt independently. The early warning level focuses on detecting changes in the maritime environment and predominantly uses autonomous platforms, satellites and long-range sensors. The monitoring level aims to classify contacts and carry out a probabilistic assessment of threats by integrating information from all available sources. The immediate protection level focuses on critical infrastructure and coordinates the operational response through the AHMDA architecture.

The operation of the AMSZ depends on **the Dynamic Threat Index (DTI)**. Whilst the DTI provides a continuous assessment of the risk level associated with each detected contact, the AMSZ uses this information to reorganise the operational space. The two concepts are complementary: the DTI answers the question **'how dangerous is the threat?'**, whilst the AMSZ answers the question **'how should the defence posture be reorganised to respond to this threat?**'.

Equally, AMSZ is inseparable from **the Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)**. Artificial intelligence may recommend changes to the configuration of the security space, but the implementation of these measures remains subject to human validation. This division of responsibilities ensures that human control over the decision-making process is maintained and allows for the integration of legal, political and strategic considerations that cannot be assessed solely through algorithms.

Applying the concept to the case of the Port of Constanța illustrates the advantages of such an approach. Instead of uniform surveillance of the entire maritime area, the system can identify in real time the sectors where the probability of an incident occurring increases and can redeploy autonomous platforms, mobile sensors and monitoring resources without altering the existing physical infrastructure. Thus, the architecture's efficiency stems not from an increase in the number of devices, but from the continuous optimisation of their use.

The AMSZ concept also offers significant advantages from the perspective of international cooperation. As risk assessment is based on the integration of information from multiple sources——the architecture can operate in a multinational environment, facilitating the exchange of data between NATO and European Union member states. In this way, the security zone is no longer limited by a state's administrative borders, but can become a distributed information structure capable of managing cross-border and multi-domain threats.

Looking ahead, the development of autonomous platforms, quantum communications and distributed artificial intelligence will transform the AMSZ into an increasingly autonomous architecture in terms of observation, analysis and the recommendation of defensive measures. However, this study argues that the legitimacy of the use of force and the assumption of legal responsibility must remain permanently within the remit of the human element. For this reason, the Adaptive Maritime Security Zone does not represent an autonomous security zone, but rather an adaptive architecture designed to support human decision-making.

In conclusion, **the Adaptive Maritime Security Zone redefines the very concept of maritime security space**. Whilst the traditional model was built around fixed geographical boundaries and a predominantly post-event response, the AMSZ introduces a dynamic, predictive and cognitive approach, in which the configuration of the defensive system continuously evolves in line with changes in the probability of threats. In this new paradigm, security is no longer the result of control over a defined maritime territory, but of the ability to constantly adapt the operational architecture to the accelerated pace of technological and strategic transformations.

4.2. Artificial intelligence and decision-making in the defence of critical maritime infrastructure

The evolution of autonomous systems and the integration of artificial intelligence into maritime security architectures are altering not only the ability to detect threats, but the very philosophy of the decision-making process. Whereas, in the past, technical systems served solely to provide information to human operators, contemporary architectures are capable of generating predictive assessments, estimating the probability of an attack occurring, and recommending, in real time, response measures tailored to the operational context. In this sense, artificial intelligence is no longer merely an analytical tool, but a multiplier of decision-making capacity.

This transformation must, however, be understood within its true limits. Artificial intelligence does not 'decide' in a legal or military sense. It reduces uncertainty and accelerates information processing, but the responsibility for the use of force and for bearing the legal consequences remains with the human operator. From this perspective, the operational value of modern systems does not stem from the complete autonomy of algorithms, but from their ability to provide the commander with the most accurate and rapid picture of the tactical situation.

In the maritime environment, a decision is always the result of a risk assessment process. Before authorising an intervention, the nature of the detected contact must be established, along with the probability that it represents a real threat, the time available before the protected target is reached, the consequences of any inaction, and the risks associated with the use of force. All these variables must be analysed simultaneously within a very short timeframe, sometimes a matter of minutes. The complexity of this process explains why modern architectures utilise algorithms capable of continuously calculating and updating the level of risk associated with each target.

Unlike traditional approaches, in which a target was classified by applying fixed rules, current systems use dynamic probabilistic models. Every piece of new information alters the existing assessment. A vessel without an AIS signal may initially appear to be a simple pleasure craft. However, if the same vessel is detected travelling at high speed towards an oil terminal, exhibits an acoustic signature characteristic of a high-powered engine, fails to respond to radio warnings and continues on a course converging with critical infrastructure, the probability that it constitutes a threat increases progressively. The decision is no longer based on a single observation, but on the successive accumulation of indicators from independent sources.

This approach allows for the introduction of an operational concept that we consider essential for the development of modern maritime defence systems: **the Dynamic Threat** *Index (DTI*). Unlike binary classifications – 'threat' or 'non-threat' – the dynamic index expresses the level of risk on a continuous scale and is automatically updated as the system receives new information. Thus, each contact within the area of interest is characterised by a value reflecting the probability of an attack occurring, and this value constantly evolves depending on the observed behaviour.

Determining such an index involves integrating several categories of factors. The technical characteristics of the platform are just one of the elements analysed. The algorithms take into account speed, acceleration, changes in direction, navigation patterns, the presence or absence of automatic identification, proximity to critical infrastructure, weather conditions, the presence of electromagnetic jamming activities, and information from external sources regarding potential

threats. By aggregating this data, the system generates a continuous risk assessment, which can be used by the commander to set operational priorities.

The advantage of such an approach lies in the fact that it allows for the simultaneous management of a very large number of contacts without overburdening human operators. Rather than every detected object requiring the same level of attention, the system automatically prioritises situations with a high potential for risk and recommends the allocation of resources based on the probable severity of the threat. In this way, the available time and resources are focused on those events that could genuinely affect the security of the protected infrastructure.

However, the use of artificial intelligence in decision-making also raises significant issues regarding the transparency of algorithms. In the defence sector, trust in an automated system cannot be built solely on the basis of statistical performance. The commander must understand the reasons why the algorithm recommends a particular course of action and be able to verify the information that led to this conclusion. For this reason, one of the major areas of development is the integration of **Explainable Artificial Intelligence (XAI)** techniques, which enable the factors influencing the system's assessment to be presented in an intelligible form.

In an architecture designed to protect the Port of Constanța, this transparency is essential. A system that recommends neutralising a platform must be able to demonstrate, in a verifiable manner, that the decision is based on the convergence of several independent sources of information and not on a single isolated observation. This reduces both the risk of false alarms and the possibility of subsequent challenges to the legality of the measures taken.

From an operational perspective, the decision-making process can be organised into successive alert levels. At the first level, the system detects and monitors the contact without initiating any active measures. At the second level, once a pre-set threshold of the dynamic threat index has been exceeded, the platform recommends intensifying surveillance and redeploying mobile sensors. In the next stage, response measures are prepared – the activation of aerial drones, the deployment of physical barriers, and the mobilisation of intervention vessels and electronic warfare systems. Only at the final level, following human validation and confirmation of the imminent nature of the threat, is the use of active countermeasures or lethal force authorised. Such a phased structure allows human control over the decision-making process to be maintained and reduces the likelihood of disproportionate reactions.

Looking ahead, the development of maritime defence systems will depend increasingly on the quality of the algorithms that manage this decision-making process. Operational superiority will no longer be determined solely by sensor performance or the number of autonomous platforms available, but by the ability of the entire architecture to transform information into a decision more quickly than the adversary. In this respect, artificial intelligence does not replace the commander, but rather extends their cognitive capacity, reducing uncertainty and providing support for the adoption of rapid, well-founded decisions that comply with the requirements of international law and contemporary military doctrines.

CHAPTER 5

Simulation of a multi-domain attack on the Port of Constanța and the response of an integrated architecture assisted by artificial intelligence

The effectiveness of a security architecture cannot be assessed solely by analysing the individual performance of its components. A system's true capability only becomes apparent when it is analysed under conditions approximating those of a real-world operation. For this reason, this chapter proposes the simulation of a complex attack scenario on the Port of Constanța, based on lessons learnt from the Black Sea conflict and current trends regarding the use of autonomous platforms, electronic warfare and integrated cyber operations.

The scenario does not aim to describe a specific historical situation, nor does it seek to attribute such an operation to a particular state or non-state actor. It represents a forward-looking

analysis designed to assess how a modern defence architecture, based on multisensory fusion and artificial intelligence, can respond to a multi-domain attack directed against critical maritime infrastructure.

It assumes a regional context characterised by heightened tensions in the Black Sea basin, an intensification of electronic warfare activities, and the emergence of intelligence regarding possible sabotage operations against energy and logistics infrastructure. In this context, the Port of Constanța represents a target of high strategic value, owing to its role in commercial transit, allied military mobility and regional energy security.

Phase I – Invisible preparation for the attack

The operation begins long before a naval drone physically appears.

In the early hours, seemingly insignificant changes are observed in the digital environment. The cyber monitoring platform detects an increase in automated attempts to scan the port's IT infrastructure; unusual fluctuations in radio communications occur in certain frequency bands; and electromagnetic spectrum sensors report intermittent emissions originating from the open sea.

Taken separately, none of these events warrants triggering an operational alert.

However, the AI platform observes that all these changes are occurring simultaneously and differ statistically from the usual behaviour of the operational environment.

The Dynamic Threat Index (DTI) begins to rise gradually.

The Adaptive Maritime Security Zone automatically adjusts its configuration.

Autonomous patrol vehicles are redeployed without human intervention to the sectors deemed most vulnerable.

VTOL aerial drones are placed on standby.

Without the adversary noticing, the system has already begun the adaptation process.

Phase II – Entry of the hostile platform into the area of interest

Approximately twenty-two nautical miles east of the Port of Constanța, one of the autonomous vehicles detects an acoustic anomaly.

The signal is weak.

Spectral analysis indicates the presence of a mechanical source consistent with a waterjet propulsion system.

Taken in isolation, the probability of classifying it as a threat is low.

However, the algorithm instantly compares this information with observations from across the entire network.

A few seconds earlier, a coastal radar had detected a very faint echo in the same area.

A smart buoy identifies the same acoustic frequency.

The system estimates that the probability of an autonomous platform being present exceeds the alert threshold.

Without transmitting continuous video streams, the autonomous vehicle activates the thermal camera locally.

The Edge AI processor analyses the images.

After just a few seconds, it identifies a thermal signature incompatible with the natural environment.

The command centre receives only the coordinates, the probabilistic classification and the compressed image.

The volume of data transmitted is several hundred times smaller than in the case of a conventional video transmission.

Reaction time is minimised.

Phase III – Fusion of information and construction of the tactical picture

There is no operator in the command centre monitoring dozens of screens simultaneously.

The AI platform integrates all available information into a single model.

The radar provides the position.

Acoustic sensors confirm the presence of a mechanical source.

The thermal camera confirms the presence of a platform.

The absence of an AIS signal rules out the possibility of a commercial vessel.

The trajectory is analysed in relation to historical traffic patterns in the area.

The system notes that the platform is avoiding commercial shipping lanes and is altering its course to reduce the time to the oil terminal.

The predictive model estimates, with a probability of over 90 per cent, that the platform's target is critical infrastructure.

The DTI exceeds the critical threshold.

The system recommends moving to the next operational level.

Phase IV – Neutralising the adversary's information chain

Before employing any kinetic measures, the proposed architecture seeks to degrade the hostile platform's command and control capabilities.

Electronic warfare systems are automatically directed towards the area of interest.

An attempt is made to interrupt radio communications and disrupt the satellite navigation channels used by the drone.

At the same time, the AI platform monitors the target's reaction.

If it loses stability or begins to perform erratic manoeuvres, the system automatically reduces the threat level and recommends continued monitoring.

However, if the platform maintains its trajectory using autonomous navigation systems, the likelihood of a fully autonomous mission increases, and the architecture moves on to the next stage.

Phase V – Integrated Intervention

At this point, all system components are operating simultaneously.

Aerial drones visually track the platform.

Autonomous maritime vehicles continue acoustic monitoring.

Radars constantly update the position.

Smart barriers near the port automatically switch to defensive mode.

Response vessels are constantly receiving updated coordinates.

The commander no longer needs to request information from each operator.

The entire operational picture is generated automatically.

The AI platform continuously calculates the time remaining until impact.

It simulates several interception scenarios.

It estimates the probability of success for each of these.

The human operator receives not only information, but also a comparative analysis of the available options.

Phase VI – Human Decision

In accordance with the **human-in-the-loop** principle, the system does not authorise the use of force.

Following multisensory confirmation of the threat and verification of all legal and operational conditions, the platform transmits a neutralisation recommendation to the commander.

The commander orders the use of the means deemed proportionate and appropriate to the situation.

Whether the intervention is carried out using electronic warfare systems, naval interception or other defensive means, the decision rests solely with human authority.

Artificial intelligence provides the necessary cognitive support, but does not replace the commander's legal and military responsibility.

Conceptual model for the dynamic assessment of maritime threats

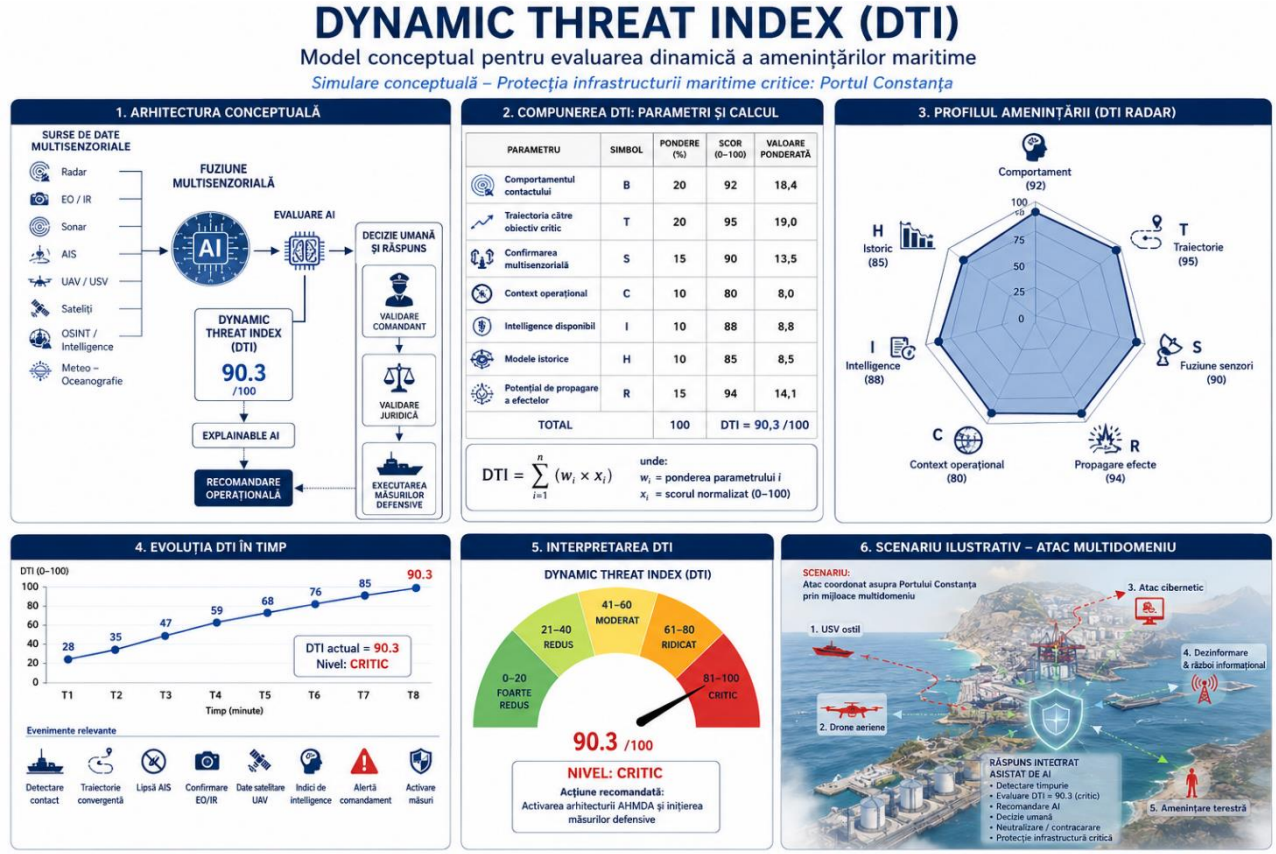

Illustrative example of a Dynamic Threat Index (DTI) calculation

Narrative example

In a conceptual simulation, an unmanned naval vehicle is detected approximately 14 nautical miles from port infrastructure. The radar system identifies a small contact travelling at a constant speed, whilst electro-optical sensors confirm the absence of markings specific to commercial shipping. At the same time, AIS analysis reveals the absence of an identification signal, and autonomous surveillance platforms transmit information regarding successive changes in the trajectory towards an area where energy infrastructure and port terminals are located.

The AHMDA architecture analysis engine integrates this information into a multi-sensor fusion process and calculates a Dynamic Threat Index of 89.1, corresponding to the critical level. Based on this result, the system recommends redeploying autonomous monitoring platforms, activating additional surveillance measures and notifying the command centre. The recommendations generated are subject to human validation before a final decision is taken, in accordance with the principles of human control and legal accountability.

Conceptual simulation – protection of critical maritime infrastructure

Parameter	Weight	Observed value	Weighted score
Contact speed	10%	85/100	8.5

Direction towards critical target	20%	95/100	19.0
Radar signature	10%	70/100	7.0
Abnormal behaviour	15%	90/100	13.5
Multisensory confirmation	15%	100/100	15.0
AI confidence level	10%	88/100	8.8
Tactical context	10%	82/100	8.2
Available intelligence	10%	91/100	9.1

Calculation

$$DTI = \sum_{i=1}^{8} (Ponder\, e_i \times Scor_i)$$

Result:
DTI = 89.1 / 100
Interpretation

DTI range	Level	Recommended action
0–20	Very low	Routine monitoring
21–40	Low	Intensified surveillance
41–60	Moderate	Multisensory confirmation
61–80	High	Command centre alert
81–100	Critical	Activation of AHMDA and initiation of defensive measures

In the illustrative example:
DTI = 89.1 → CRITICAL level

Lessons learnt from the simulation

The scenario analysed highlights the fact that the effectiveness of a modern defence architecture is not determined by the individual performance of a platform or sensor, but by the synchronisation of the entire information ecosystem. Autonomous vehicles, smart radars, acoustic sensors, electro-optical cameras, resilient communications and AI-assisted command platforms all contribute to building a common tactical picture, reducing uncertainty and extending the time window available for reaction.

The analysis also confirms that the success of defence does not depend exclusively on the physical neutralisation of the hostile platform. In many situations, early detection, the disruption of the adversary's communications, the dynamic redeployment of sensors and the continuous adaptation of the defensive posture can prevent the adversary from achieving their objective without the need for the immediate use of force. In this regard, artificial intelligence should not be viewed merely as a tool for automation, but as a central element of operational resilience and information superiority.

5.2. Artificial intelligence, decision-making autonomy and the limits on the use of force in the protection of critical maritime infrastructure

The development of autonomous surveillance and defence systems is profoundly altering the relationship between technology and the military decision-making process. Whilst until recently IT systems played a predominantly passive role, limited to collecting and displaying information,

the new generations of artificial intelligence-assisted platforms are capable of analysing very large volumes of data, identifying operational patterns, anticipating the evolution of a threat and formulating recommendations on response measures. However, this development raises a fundamental question: to what extent can the role of algorithms be extended in a field where decisions can have lethal consequences and international legal implications?

The issue is not exclusively technological. Essentially, it concerns the distribution of responsibility between automated systems and the human factor. The more effective artificial intelligence becomes at detecting and classifying threats, the greater the temptation to extend its autonomy to subsequent stages of the decision-making process. However, under contemporary international law and in the military doctrines of democratic states, the use of force cannot be reduced to the result of an algorithmic calculation, regardless of its level of accuracy.

In the case of the protection of critical maritime infrastructure, this distinction is particularly important. An autonomous platform approaching an oil terminal could be a civilian vessel in distress, a drone used for research, a commercial vessel without a functioning AIS system, or a platform intended for a deliberate attack. Even if artificial intelligence estimates a very high probability of a threat, this assessment cannot replace the legal and operational analysis that the commander is obliged to carry out before authorising the use of force.

This obligation derives from the general principles of international law and from the rules governing the use of force by state authorities. Any neutralising measure must comply with the criteria of necessity, proportionality and precaution. Necessity requires that the intervention be indispensable for the protection of a legitimate interest and that there be no alternative means capable of removing the threat. Proportionality requires that the intensity of the response be commensurate with the level of risk, whilst avoiding excessive effects on persons, property or the environment. The precautionary principle obliges the competent authority to verify, as far as possible, the true nature of the target before authorising the use of force.

Artificial intelligence can support the application of these principles, but it cannot replace them. On the contrary, the integration of algorithms into the decision-making process imposes additional standards regarding the verifiability and transparency of the conclusions generated. In practice, this means that every recommendation made by the system must be accompanied by a justification that is comprehensible to the human operator: which sensors contributed to the classification, what the confidence level is for each observation, which alternative hypotheses have been ruled out, and what the limitations of the automated assessment are. In the absence of this information, the commander cannot exercise effective control over the decision-making process and cannot genuinely assume legal responsibility for the measures ordered.

From this perspective, the concept of **Explainable Artificial Intelligence (XAI)** takes on strategic importance. In the military context, explainability is not merely a technical requirement concerning the interpretability of algorithms, but a condition for the legality of the decision-making process. A system that generates recommendations without being able to explain the reasons behind them risks reducing the human operator to a mere formal validator of a decision already adopted by the algorithm, which would contravene the very rationale of *the 'human-in-the-loop'* principle.

At the same time, the opposite tendency—namely, over-reliance on human judgement to the detriment of the results produced by intelligent systems—must also be avoided. Research in the field of ' ' of human factors demonstrates that operators are exposed both **to the phenomenon** of **'automation** bias' – characterised by the uncritical acceptance of recommendations generated by automated systems – and to **the** phenomenon of **'algorithm aversion**', manifested by the systematic rejection of algorithmic conclusions even when these are more accurate than human assessment. An effective architecture must simultaneously mitigate both risks, ensuring a balance between the operator's expertise and the analytical capabilities of artificial intelligence.

In the case of the protection of the Port of Constanța, this approach involves organising the decision-making process across several successive levels. Algorithms identify and classify the threat; the command and control platform aggregates information from all available sources; and the commander validates the proposed measures only after analysing the operational context and the

legal implications of the intervention. In this way, artificial intelligence accelerates the decision-making cycle without removing human control over critical stages.

An additional aspect, which has not been sufficiently analysed in the specialist literature, concerns the obligation to preserve digital evidence. In the event of an incident, the authorities must be able to reconstruct the entire decision-making process: the data received from sensors, the recommendations made by algorithms, the operators' interventions and the final order to use force. For this reason, modern architectures must include mechanisms for *secure* logging, precise time synchronisation and protection against subsequent data tampering. These records are not only useful for post-incident analysis, but also for establishing potential legal liability, for operational auditing and for the continuous improvement of the algorithms used.

From a theoretical perspective, the development of autonomous systems also requires a re-examination of the concept of operational responsibility. In a distributed architecture, where information is generated by different sensors, analysed by multiple algorithms and validated by operators in separate locations, responsibility can no longer be reduced to the actions of a single individual. It must be understood as the result of a complex decision-making chain, within which each component fulfils a precise and verifiable function. This necessitates the definition of clear procedures regarding operators' responsibilities, the limits of system autonomy, and the criteria for validating recommendations generated by artificial intelligence.

In conclusion, the integration of artificial intelligence into modern maritime security architectures does not diminish the importance of the human factor, but rather redefines its role. The captain is no longer the primary processor of information, but the primary guarantor of the legality and legitimacy of the decision. Artificial intelligence provides operational knowledge, speeds up analysis and reduces uncertainty; however, the use of force remains a human decision, based on a legal, military and ethical assessment of the specific situation. It is precisely this complementarity between the cognitive capacity of algorithms and the responsibility of human authority that forms the foundation of a modern, efficient architecture that complies with the requirements of international law.

CHAPTER 6

Adaptive, Human-Centred Maritime Decision-Making Architecture (AHMDA): a conceptual model for the protection of critical maritime infrastructure

Figura 1

The rapid transformation of the maritime security environment demonstrates that the development of more advanced sensors or the integration of artificial intelligence into surveillance processes is not, in itself, a sufficient solution to address contemporary threats. Technological developments in recent years have led to the emergence of an operational paradox. On the one hand, modern systems generate unprecedented volumes of information from radars, electro-optical sensors, autonomous platforms, satellite systems and digital infrastructure. On the other hand, the exponential growth in the volume of data does not automatically lead to improved decision-making. On the contrary, in the absence of an architecture capable of organising, interpreting and prioritising this information, technological superiority risks being negated by operators' cognitive overload and delayed reactions in an operational environment characterised by speed and uncertainty.

This reality highlights the need for a paradigm shift. The central issue in maritime security is no longer the detection of a threat, but the rapid and responsible transformation of information into a legitimate operational decision. In this context, this study proposes the concept **of the Adaptive Maritime Decision-Making Architecture, centred on the human** as an integrator, as a model for the design of modern defence systems for critical maritime infrastructure.

AHMDA can be defined as an adaptive command and control architecture in which artificial intelligence, autonomous systems and multisensory fusion are used for the detection, classification and predictive assessment of threats, whilst the authority to use force, amend the rules of engagement and assume legal responsibility remains permanently with the human element. The model does not aim to replace human decision-making, but rather to redistribute cognitive functions between algorithms and operators, so that each component contributes in the area where it offers the highest level of performance.

The originality of this architecture lies in the fact that it does not treat artificial intelligence as an autonomous tool, but as a constituent element of an information ecosystem in which technology and human expertise complement one another. Algorithms are used to process very

large volumes of data simultaneously, to identify relationships between seemingly independent observations, and to estimate the likely evolution of the tactical situation. The human factor comes into play where contextual judgement, legal assessment, strategic interpretation and accountability for the consequences of the decision taken are required. In this configuration, artificial intelligence does not replace the commander, but rather enhances their cognitive capacity, reducing the time required for analysis without compromising human control over the critical stages of the decision-making process.

The AHMDA model is based on the premise that the decision-making process cannot be understood as a linear sequence of activities, but rather as an adaptive system characterised by continuous loops of observation, interpretation and recalibration. Any new information alters the existing assessment and may lead to the reconfiguration of the entire operational architecture. In this sense, surveillance, analysis and response are not distinct stages, but interdependent components of a continuous process of adaptation to the dynamics of the security environment. The proposed architecture does not react solely to events as they occur, but seeks to anticipate the evolution of threats and to adapt the allocation of available resources proactively.

A key feature of the model is the integration of the concept of multisensory convergence. Rather than placing absolute trust in information from a single sensor, the architecture builds a level of certainty by continuously correlating radar, electro-optical, thermal, acoustic, satellite and cyber data. Each observation contributes to the probabilistic assessment of the situation, and the system's confidence level increases as independent sources confirm the same operational hypothesis. In this way, the risk of false alarms is reduced, and critical decisions are based on an operational picture consolidated through the convergence of multiple categories of information.

A second defining feature is the adaptability of the architecture. Unlike traditional systems, designed to operate on the basis of fixed configurations, AHMDA involves the continuous redeployment of sensors, autonomous platforms and response resources in response to changes in the operational environment. Thus, the concept of **the Adaptive Maritime Security Zone (AMSZ)**, introduced in previous chapters, becomes the spatial expression of this adaptability. The security zone is no longer a pre-defined perimeter, but a dynamic configuration, continuously recalculated on the basis of predictive risk assessment. Depending on how the tactical situation develops, the architecture can expand, contract or redeploy resources without waiting for a manifest threat to emerge.

Following the same logic, the risk assessment process is supported by **the Dynamic Threat Index (DTI)**, designed as a mechanism for the continuous assessment of the threat level associated with each detected target. Unlike the binary classifications typical of traditional systems, the DTI reflects the dynamic nature of the operational environment and allows the assessment to be constantly updated in line with newly collected information. This index does not automatically determine the response measure, but provides the commander with a reasoned assessment of the probabilistic evolution of the situation, helping to prioritise resources and organise the operational response.

Figure 1. Adaptive, human-centred maritime decision-making architecture (AHMDA)

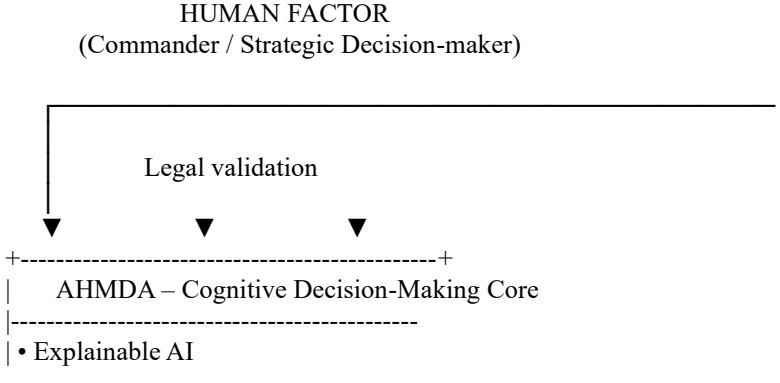

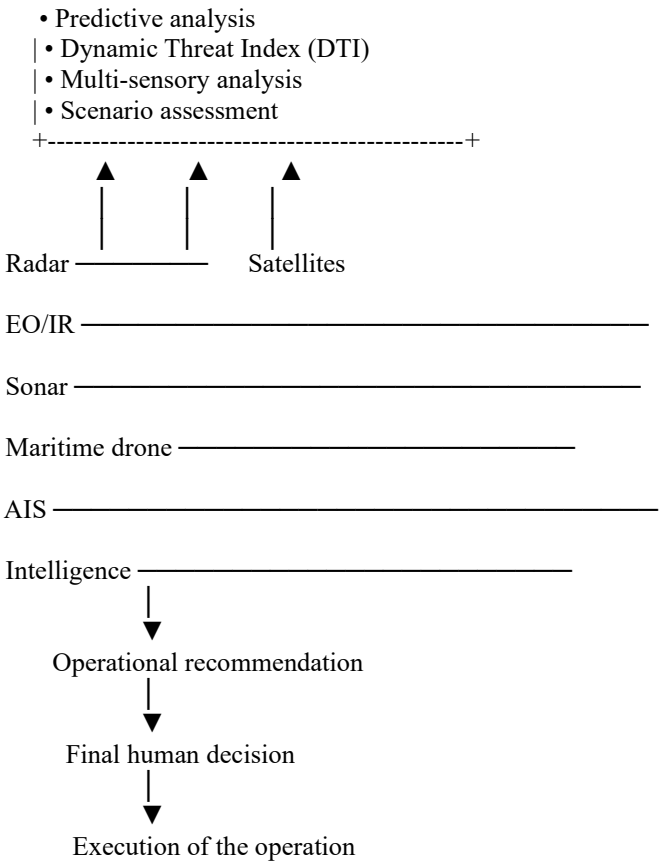

A fundamental aspect of the AHMDA architecture concerns the transparency of the decision-making process. As algorithms become more sophisticated, there is a growing risk that the recommendations generated will be perceived as the results of opaque processes that are difficult to understand and verify. In a field where decisions may involve the use of force and have significant legal implications, this opacity is incompatible with the requirements of institutional accountability. For this reason, the proposed model incorporates the principles of Explainable Artificial Intelligence (*EAI*), on the basis that every algorithmic recommendation must be justifiable by indicating the sources used, the confidence level associated with each observation, and the reasoning that led to the conclusion reached. Explainability is not merely a technical feature, but a prerequisite for the exercise of effective human control over the decision-making process.

At the same time, AHMDA attaches particular importance to the resilience of the information architecture. Recent conflicts demonstrate that cyber-attacks and electronic warfare operations are primarily aimed at disrupting information flows and undermining the decision-making process. Consequently, the proposed architecture requires the existence of distributed communications networks, edge processing mechanisms (*Edge AI*) and degraded-mode operational capabilities, so that the temporary loss of certain sensors or communications links does not lead to the collapse of the entire system. Resilience is not viewed solely as a technical characteristic, but as a systemic property, indispensable to maintaining the continuity of the decision-making process.

Perhaps the most important feature of the AHMDA model, however, is the reaffirmation of the central role of the human factor. In contemporary literature, there is a tendency to evaluate the performance of autonomous systems in terms of the degree of automation they achieve. The proposed architecture adopts a different perspective. The level of performance is not determined by a reduction in human intervention, but by the quality of collaboration between the operator and the algorithm. Artificial intelligence is used to reduce uncertainty, accelerate analysis and expand the commander's cognitive capacity, without it becoming the decision-maker. The authority to use force, the interpretation of the strategic context and the assumption of legal responsibility remain permanently within the remit of the human factor. In this way, the *'human-in-the-loop'* principle is

transcended in a functional sense: the human is not merely a formal validator of algorithmic recommendations, but the centre of gravity of the entire decision-making architecture.

From a doctrinal perspective, AHMDA proposes a shift in focus in the design of maritime security systems. The objective is no longer the development of increasingly sophisticated autonomous platforms, but rather the construction of an information ecosystem in which technology, predictive analysis, resilient communications and human expertise function as elements of a unified mechanism. Operational superiority stems from the quality of the relationships between these components and from the architecture's ability to rapidly transform information into well-founded, proportionate decisions that comply with the requirements of international law.

From this perspective, the Adaptive Human-Centric Maritime Decision Architecture is not merely a technological model, but a conceptual framework designed for the development of future systems to protect critical maritime infrastructure. By integrating artificial intelligence into an architecture centred on human responsibility, adaptability and resilience, the model offers a direction for development that is compatible both with the operational requirements of contemporary conflicts and with the legal and ethical requirements governing the use of force in the maritime domain. In this regard, AHMDA can serve not only as an analytical tool for evaluating existing architectures, but also as a benchmark for designing the next generation of command and control systems for maritime security.

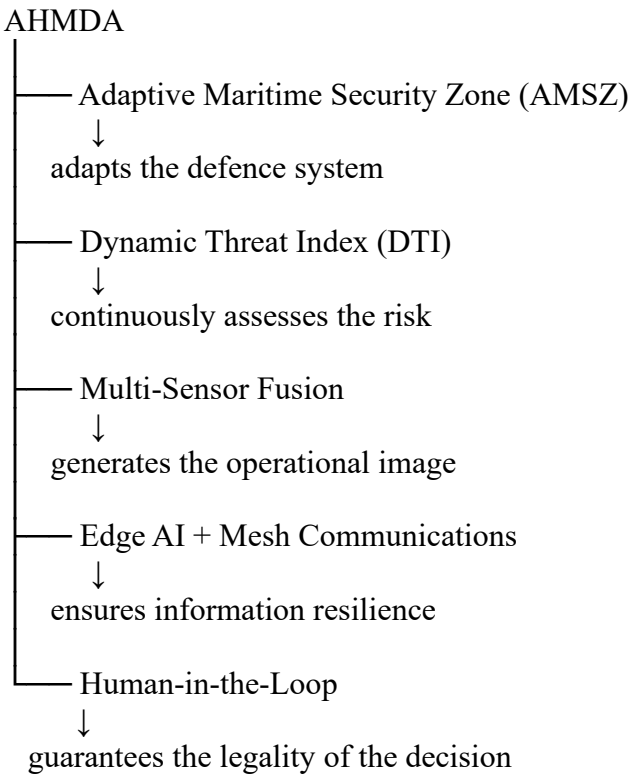

6.1. Conceptual validation of the Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA) model

Any conceptual model intended for the development of a new operational architecture must be assessed not only in terms of its theoretical coherence, but also in terms of its ability to address the limitations identified in existing systems. In the case of AHMDA, the objective is not to replace current maritime security doctrines, nor to substitute the command and control architectures already in use by NATO member states or the European Union. The proposed model aims to integrate

recent advances in the fields of artificial intelligence, multi-sensor fusion and autonomous systems into a unified conceptual framework capable of addressing the specific characteristics of contemporary threats in the maritime domain.

The need for such an architecture stems from the profound change in the nature of conflicts. Whereas, in the past, maritime threats were primarily posed by conventional naval platforms—which were relatively easy to identify and track—the current operational environment is characterised by the proliferation of autonomous vehicles, hybrid operations, cyber-attacks and electronic warfare. These developments significantly reduce the time available to react and increase the degree of uncertainty in the process of identifying threats. In this context, operational superiority is no longer determined exclusively by the individual performance of sensors or combat platforms, but by the speed and quality of the process through which information is transformed into a decision.

Traditional command and control architectures were developed at a time when information flows were far less complex. They rely on a command centre that collects data from sensors, analyses it and transmits orders to operational units. This model remains effective in many situations, but encounters difficulties when the number of data points increases exponentially, and information arrives simultaneously from a wide variety of sources and must be analysed within an extremely short timeframe. Under such conditions, the limiting factor is no longer the technical performance of the sensors, but the cognitive capacity of the operators and the time required to integrate and interpret the available information.

The AHMDA model addresses this challenge by redistributing cognitive functions between artificial intelligence and the human factor. Algorithms take over repetitive tasks, characterised by the processing of very large volumes of data, pattern recognition and probabilistic risk assessment. Human operators focus their efforts on stages where professional experience, contextual judgement and legal responsibility are indispensable. In this configuration, artificial intelligence does not diminish the commander's role, but rather enables them to make more efficient use of their available time and cognitive resources.

A key advantage of the proposed architecture lies in reducing the vulnerabilities caused by information overload. In conventional systems, an increase in the number of sensors almost inevitably leads to an increase in the volume of information that operators must analyse. In AHMDA, this relationship is reversed. The introduction of additional sensors does not result in a proportional increase in cognitive load, as the information is filtered, correlated and prioritised before reaching the decision-maker. The commander no longer receives all the data collected by the system, but only that information which is operationally relevant and requires human intervention.

Conceptual validation of the model can also be achieved by reference to the principles of resilience. Recent conflicts demonstrate that the initial actions taken against critical infrastructure frequently aim to disrupt communications, compromise IT systems and disrupt information flows. In a centralised architecture, such actions can lead to the disruption of the entire decision-making process. In contrast, AHMDA utilises distributed networks, edge computing and adaptive resource reallocation mechanisms, which enable functionality to be maintained even in the event of the loss of individual components. Resilience does not stem from the invulnerability of each individual element of the system, but from the ability of the entire architecture to continue the decision-making process under degraded conditions.

From a legal perspective, the proposed model aims to reconcile two requirements that are often presented as antagonistic: accelerating the decision-making process and maintaining human control over the use of force. In the specialist literature, there is a tendency to consider that an increase in the autonomy of systems inevitably entails a reduction in the role of the human factor. AHMDA proposes a different approach. Autonomy is used for information processing and for evaluating available alternatives, but the decision on the use of force remains the responsibility of the commander. This functional distinction allows the benefits offered by artificial intelligence to be harnessed without undermining the principles of individual accountability and effective human control.

The model also offers a significant advantage in terms of interoperability. As it does not involve replacing existing infrastructure but rather integrating it into a shared information ecosystem, AHMDA can be implemented gradually. Existing radar systems, autonomous platforms, coastal surveillance infrastructure and command centres can be connected via multi-sensor fusion mechanisms and AI-assisted analysis platforms, without requiring a complete overhaul of the national maritime security architecture. For Romania, this feature is particularly important, as it allows the capitalisation on investments made in the SCOMAR system and other maritime capabilities, complementing them with emerging technologies and new analytical mechanisms.

From the perspective of future development, AHMDA can also serve as a methodological framework for evaluating other maritime security systems. The concept is not limited to the protection of the Port of Constanța, nor to the specific characteristics of the Black Sea basin. Its principles – multisensory convergence, adaptability, explainability, resilience and the centrality of the human factor – can be used to analyse and design architectures intended to protect commercial ports, offshore energy terminals, submarine cables, offshore wind farms and other critical infrastructure exposed to threats posed by autonomous platforms.

From this perspective, AHMDA should not be understood as a closed technological solution or as a product intended for a specific category of users. It represents an open conceptual model, capable of being adapted in line with developments in technology, military doctrines and applicable legal norms. Its value does not derive from the technical specifications of a particular platform, but from its ability to organise the relationship between information, artificial intelligence and human responsibility in a coherent manner that is compatible with the operational requirements of contemporary maritime security.

CHAPTER 7

Implications for the maritime security architectures of NATO and the European Union

The transformation of the security environment in the Black Sea basin, the intensification of hybrid operations and the proliferation of autonomous systems are prompting a profound reassessment of how international organisations understand the protection of critical maritime infrastructure. Whilst, in the past, the maritime security architectures developed by NATO and the European Union were primarily aimed at monitoring trade routes, combating piracy, controlling maritime borders and safeguarding freedom of navigation, recent technological developments have significantly broadened the scope of these concerns. Today, offshore energy infrastructure, submarine communications cables, LNG terminals, commercial ports and the digital ecosystems that support their operation are regarded as essential components of collective security.

This shift in perspective is not merely an adaptation to new threats, but reflects the very transformation of the concept of maritime power. The contemporary maritime domain is no longer defined exclusively by its geographical extent, but by the interdependence between physical infrastructure, digital networks and information flows that underpin the functioning of the global economy. In these circumstances, the protection of critical maritime infrastructure requires the development of architectures capable of simultaneously integrating naval, air, cyber, space and information dimensions into a unified command and control process.

Within NATO, this development is reflected in the emergence of concepts such as **Multi-Domain Operations (MDO)** and **Joint All-Domain Command and Control (JADC2)**, which aim to integrate information from all operational domains into a common picture, available to decision-makers in real time. The fundamental idea underpinning these initiatives is that strategic advantage no longer derives from superiority in a single operational d , but from the ability to rapidly correlate information from multiple sources and coordinate the response in an integrated manner. In this context, the maritime environment is no longer treated as an isolated operational theatre but becomes one of the components of a multi-domain information ecosystem.

At the same time, NATO has stepped up its focus on the protection of critical undersea infrastructure, particularly following the incidents involving energy pipelines and undersea cables in recent years. The establishment **of the Maritime Centre for the Security of Critical Undersea Infrastructure** reflects the recognition that maritime infrastructure is no longer merely an economic asset, but an indispensable element for the functioning of modern societies and for the Alliance's ability to respond. Protecting such infrastructure requires the development of systems capable of rapidly identifying anomalous activity, integrating information from a wide range of sources, and generating early warnings regarding the evolution of threats.

The European Union is moving in a similar direction. **The EU Maritime Security Strategy**, updated in 2023, places particular emphasis on strengthening the resilience of critical infrastructure, increasing interoperability between civil and military authorities, and developing a common maritime picture through the continuous exchange of information. Instruments such as **the Common Information Sharing Environment (CISE)**, the work **of the European Maritime Safety Agency (EMSA)** and operational cooperation with **Frontex** aim to build a distributed information infrastructure, in which data collected by different authorities is integrated into a common operational picture.

When analysed comparatively, the initiatives of NATO and the European Union highlight the existence of strategic convergence. Both organisations prioritise information integration, the development of interoperability and the use of digital technologies to speed up the decision-making process. However, existing strategic documents generally treat artificial intelligence as a supporting technology, without proposing a conceptual architecture that explicitly defines the relationship between algorithms, autonomous systems and human accountability in the decision-making process.

In this regard**,** we consider that the **Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)** model can complement these initiatives without replacing them. Unlike predominantly technological approaches, AHMDA proposes a framework for organising the decision-making process in which the integration of artificial intelligence is inseparable from the principles of human responsibility, algorithmic explainability and the legitimacy of the use of force. The model does not introduce a new command and control system competing with existing architectures, but rather provides an additional level of conceptualisation regarding the distribution of cognitive functions between human operators and intelligent systems.

This complementarity is also evident in relation to the development of multi-domain operations. Whilst JADC2 aims to connect platforms and sensors within a common information network, AHMDA focuses on how this information is transformed into a legitimate operational decision. From this perspective, the two models are not mutually exclusive, but operate at different levels of the same architecture: JADC2 addresses the question **of how information flows**, whilst AHMDA addresses the question **of how this information is used to make a decision that complies with operational and legal requirements**.

Similarly, the concept of **the Adaptive Maritime Security Zone (AMSZ)** can be interpreted as an operational extension of the Common Maritime Picture developed by NATO and the European Union. Whilst current systems aim to build as comprehensive a picture of the maritime situation as possible, the AMSZ introduces an additional, adaptive dimension, whereby the deployment of sensors and autonomous platforms is continuously adjusted based on predictive risk assessment. Thus, the operational picture no longer serves merely a descriptive role, but becomes an active tool for reorganising the defensive posture.

At the same time, **the Dynamic Threat Index (DTI)** provides a methodology for prioritising detected contacts, facilitating the management of a very large number of events in an environment characterised by information overload. In the context of NATO operations and European cooperation, such a mechanism could help standardise the way in which different authorities assess and classify maritime threats, reducing differences in assessment between national and multinational structures.

From Romania's perspective, these developments are of particular strategic importance. Its location on the eastern border of NATO and the European Union gives the Black Sea a vital role in the Euro-Atlantic security architecture. In this context, the development of conceptual models compatible with the strategic directions of the two organisations is not merely an academic exercise, but a prerequisite for strengthening interoperability and enhancing the resilience of national maritime infrastructure.

In conclusion, the analysis demonstrates that current trends within NATO and the European Union are converging towards the development of integrated information architectures, based on the continuous exchange of data and the use of digital technologies. However, the integration of artificial intelligence raises new challenges regarding the organisation of the decision-making process, the distribution of responsibility and the maintenance of human control over the use of force. Through the concept **of Adaptive Human-Centric Maritime Decision Architecture**, this study proposes a conceptual framework that complements these developments and offers a possible direction for the future maritime security architectures of NATO and the European Union, in which technological superiority is matched by legal legitimacy, transparency and operational resilience.

CHAPTER 8

Artificial intelligence and states' obligations regarding the protection of critical maritime infrastructure: towards a new dimension of the duty of care

Figura 6

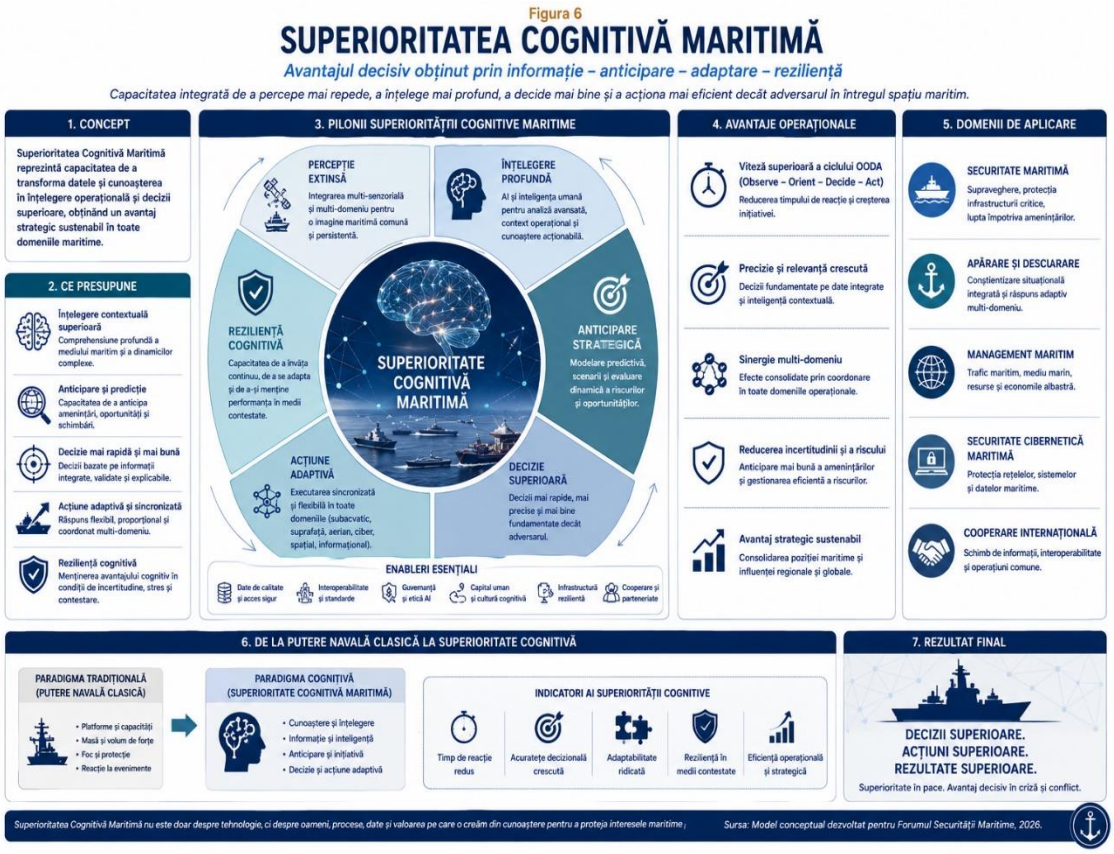

The protection of critical maritime infrastructure can no longer be analysed solely from the perspective of technological development or the modernisation of military capabilities. The integration of artificial intelligence into surveillance, command and control architectures has legal implications that extend beyond the internal organisation of states and influence the way in which they fulfil their international obligations regarding the prevention of threats, the protection of essential infrastructure and the maintenance of the ‘ ’ of maritime security. In this context, the use of artificial intelligence is not merely a technological option, but is progressively becoming a relevant

factor in assessing a state's conduct in relation to the standards of due diligence required by international law.

Traditionally, the duty of due diligence has been understood as the state's obligation to take all reasonable measures at its disposal to prevent harm to other states or to interests protected by international law. The nature of this obligation is one of conduct rather than of result. International law does not guarantee that no incident will occur, but it imposes on states the obligation to organise and exercise public authority in a reasonable and effective manner to prevent foreseeable risks.

This obligation has undergone constant development in international case law and in the practice of international organisations. It is found in the fields of environmental protection, the law of the sea, counter-terrorism, cyber security and the protection of human rights, and has been progressively adapted to the emergence of new categories of risk. From this perspective, the proliferation of autonomous platforms and the use of artificial intelligence in the maritime environment raises the question of whether the standard of due diligence can remain unchanged in a context where technology enables the detection and anticipation of threats that previously could not be identified.

This question is particularly relevant to critical maritime infrastructure. Commercial ports, oil terminals, offshore installations, submarine cables and energy pipelines are essential to the functioning of contemporary economies and to regional security. Attacks against these targets can have economic, humanitarian and environmental consequences that extend far beyond the territory of the affected state. Consequently, the obligation to protect cannot be interpreted solely as a domestic responsibility, but must also be analysed in the light of the international community's interest in maintaining the security of maritime infrastructure.

In this new technological reality, the assessment of a state's conduct can no longer be based solely on the existence of conventional surveillance methods. If technologies based on artificial intelligence enable the early detection of threats, the reduction of false alarms and the optimisation of decision-making, the question arises as to whether the failure to utilise such capabilities might influence the assessment of the reasonableness of the measures adopted by the state. It cannot be argued that there is a general obligation to implement any new technology. However, as certain solutions become mature, accessible and integrated into the operational practices of a significant number of states, the standard of due diligence may gradually change.

This development is characteristic of international law. The content of due diligence obligations is not static, but constantly adapts in line with advances in scientific knowledge, technology and international practices. In the field of environmental protection, for example, the development of new monitoring methods has influenced the assessment of the reasonable measures that states are obliged to adopt. A similar process can also be observed in the field of maritime security, where artificial intelligence is beginning to redefine what may be considered a reasonable level of risk prevention and anticipation.

From this perspective, this study proposes extending the analysis of the duty of care by introducing the concept of **'cognitive due diligence**'. This concept does not aim to create a new, autonomous legal obligation, but rather to describe an emerging dimension of the traditional duty of care. It expresses the idea that states must use, in a reasonable and proportionate manner, the cognitive and technological tools available to them for the early identification of threats that may affect critical maritime infrastructure. To the extent that artificial intelligence enables a significant reduction in risks through anticipation and predictive analysis, the systematic disregard of such capabilities may become a relevant factor in assessing a state's conduct.

The concept of *cognitive due diligence* does not entail an obligation to implement a specific software platform or algorithm. Due diligence continues to be assessed in accordance with each state's actual capabilities, available resources and the specific operational context. However, it does entail an obligation to establish institutional frameworks, to carry out ongoing risk assessments and to progressively adapt protective mechanisms in line with technological developments. In this

regard, the obligation does not concern the technology itself, but rather the authorities' ability to make reasonable use of the available tools to prevent foreseeable harm.

Applying this concept to the field of critical maritime infrastructure leads to a significant shift in perspective. The state's responsibility can no longer be assessed solely in terms of its response after an incident has occurred. Equally relevant is the way in which it organises risk anticipation, integrates information from multiple sources and utilises artificial intelligence to reduce vulnerabilities before the threat materialises. In this context, architectures such as **the Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)** represent not only technological solutions but also institutional mechanisms through which the state can fulfil its prevention obligations more effectively.

A further dimension concerns the obligation of international cooperation. Threats to maritime infrastructure are rarely confined to the jurisdiction of a single state. Autonomous vehicles can traverse multiple maritime zones, cyber-attacks are carried out via distributed infrastructure, and relevant information is held simultaneously by civil and military authorities and international organisations. In these circumstances, due diligence also entails the development of effective mechanisms for information-sharing, interoperability and coordination. Artificial intelligence can enhance the effectiveness of these mechanisms, but it cannot replace the legal obligation of states to cooperate in preventing common threats.

The same reasoning applies to the obligation to investigate. In the event of an incident, the use of artificial intelligence-assisted systems generates very large volumes of digital evidence regarding the sequence of events and the decision-making process. Preserving this information, ensuring its integrity and being able to reconstruct the chronology of the incident become essential components of an effective investigation. Thus, the duty to investigate takes on a new technological dimension, in which the auditability of algorithms and the traceability of data are just as important as traditional physical evidence.

Looking ahead, the evolution of artificial intelligence will inevitably influence the international standard of reasonable conduct. As cognitive architectures become an integral part of modern maritime security systems, it is likely that the assessment of the duty of care will, to an ever-greater extent, include an analysis of how states utilise these capabilities to anticipate and prevent threats. This process does not automatically entail the emergence of new obligations, but rather the reinterpretation of existing obligations in the light of technological progress and new possibilities for prevention.

In conclusion, the integration of artificial intelligence into the protection of critical maritime infrastructure does not alter the fundamental nature of states' international obligations, but it does influence the practical manner in which these are fulfilled and assessed. Diligence, prevention, cooperation and investigation remain the legal pillars of maritime security, but their content evolves alongside the development of technologies capable of reducing uncertainty and increasing the efficiency of the decision-making process. Within this new paradigm, the responsible use of artificial intelligence is not merely an operational advantage, but becomes an expression of a state's capacity to fulfil its international obligations in a security environment undergoing constant transformation.

8.1. The Maritime AI Readiness Maturity Model (MARMM): a model for assessing the maturity of artificial intelligence implementation in maritime security

Figura 4

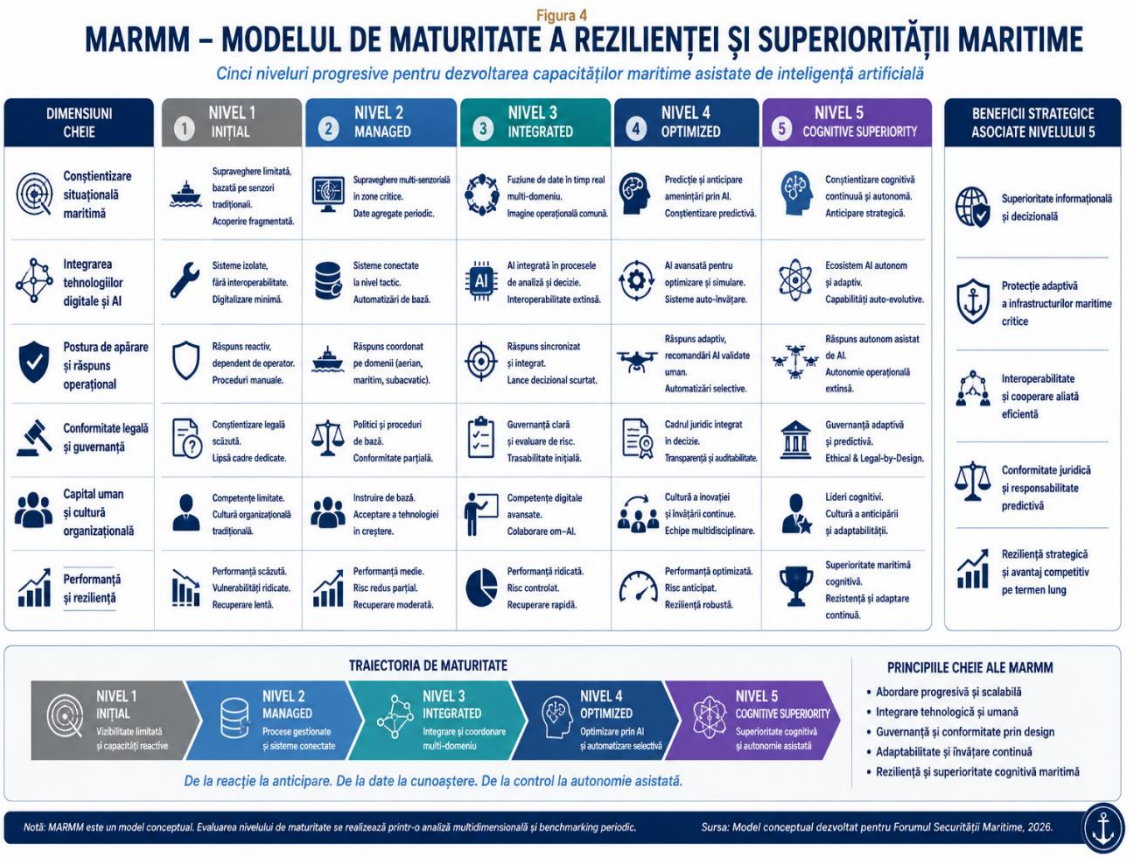

The digital transformation of maritime security is often analysed in relation to individual technologies, such as autonomous systems, artificial intelligence, distributed communications or smart sensors. Whilst such an approach is useful for describing technological progress, it does not allow for an assessment of the actual level of readiness of an organisation or a state to integrate these capabilities into a coherent operational architecture. In practice, the existence of autonomous platforms or high-performance algorithms does not automatically guarantee increased operational efficiency if they are not integrated into a decision-making process adapted to the new technological realities.

At present, there is no single model that allows for the assessment of the maturity of artificial intelligence implementation in the field of maritime security. Most assessments focus on the level of digitalisation, the performance of equipment or the degree of automation, without simultaneously analysing interoperability, the decision-making process, legal governance and systemic resilience. Consequently, this study proposes **the Maritime AI Readiness Assessment Model (MARMM)**, a conceptual tool designed to progressively assess the integration of artificial intelligence into architectures for the protection of critical maritime infrastructure.

The model is based on the premise that the maturity of an architecture is not determined solely by the complexity of the technology used, but by its ability to transform information into decisions in a secure, interoperable manner that complies with the state's legal obligations. From this perspective, maturity is the result of the convergence of technology, institutional organisation, decision-making and governance.

Level I – Conventional Surveillance

The first level corresponds to traditional maritime security architectures, based on independent sensors and predominantly manual analysis of information. Radars, surveillance cameras and AIS systems operate in parallel, and data integration is carried out by human operators. Platforms exchange information to a limited extent, and decision-making relies heavily on staff

experience and standardised procedures. The capacity for anticipatory is limited, and a response is triggered mainly following the explicit identification of a threat.

This level characterises most systems developed over the last two decades and represents the starting point for digital transformation.

Level II – Multisensory integration

The second level involves the existence of a common information fusion platform. Data from radars, optical sensors, thermal cameras, sonars and other sources are automatically correlated to construct a unified operational picture. The operator no longer interprets each information stream independently, but receives an integrated representation of the tactical situation.

At this stage, the first significant operational advantage emerges: a reduction in false alarms and an increase in the speed of identifying relevant contacts. However, the analysis remains predominantly descriptive, and the system does not yet generate complex predictive assessments.

Level III – Artificial intelligence for detection and predictive analysis

The third level marks the effective integration of artificial intelligence algorithms into the surveillance process. Systems are not limited to aggregating information, but automatically identify anomalies, classify detected objects and estimate the probability of incidents occurring based on machine learning models.

At this stage, concepts such as distributed processing (*Edge AI*), behavioural analysis and intelligent data fusion emerge. Autonomous platforms become active nodes in the information architecture, and risk assessment begins to be carried out continuously.

This level corresponds to the introduction **of the Dynamic Threat Index (DTI)** and represents the transition from reactive surveillance to threat anticipation.

Level IV – Cognitive Command and Control Architecture

The fourth level is characterised by the full integration of artificial intelligence into the decision-making cycle. Algorithms not only identify threats, but also construct alternative response scenarios, assess the consequences of each option and recommend the optimal allocation of available resources.

At this stage, adaptive architectures such as **the Adaptive Maritime Security Zone (AMSZ)** emerge, and autonomous platforms automatically adjust their position and mission based on a continuous assessment of the operational environment. The decision-making process is supported by predictive models, but the use of force remains subject to human validation.

The defining characteristic of this level is the transformation of the command centre into a cognitive system capable of simultaneously managing a very large number of contacts and operational scenarios.

Level V – Adaptive Human-Centric Maritime Decision Architecture

The top level of the MARMM model is represented by the full implementation of the **Adaptive Human-Centric Maritime Decision Architecture (AHMDA)** concept. At this stage, all components of the architecture function as a unified information ecosystem, in which artificial intelligence, autonomous platforms, resilient communications and human operators collaborate continuously to reduce uncertainty and optimise the decision-making process.

What sets this level apart is not the high degree of automation, but the integration of legal governance into the technological architecture. The explainability of algorithms, the auditability of the decision-making process, effective human oversight and compliance with the principles of international law become intrinsic components of the system, rather than mere external requirements.

Thus, maximum maturity is not defined by the complete autonomy of the technology, but by the balance between algorithmic performance and institutional accountability.

Applicability of the MARMM model

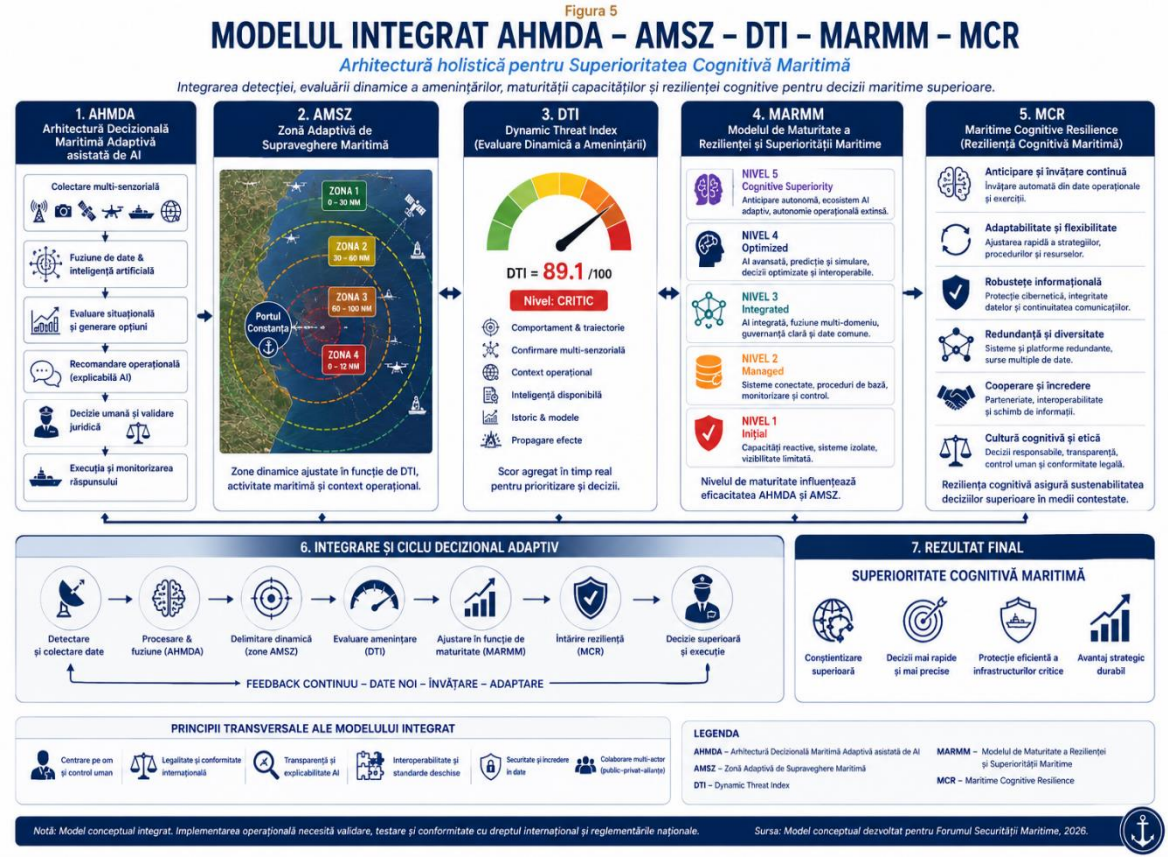

The proposed model can be used for several purposes. Firstly, it allows for a comparative assessment of the level of development of various national maritime security architectures. Secondly, it can serve as a tool for investment planning and establishing modernisation stages. Thirdly, it provides a common analytical framework for cooperation between NATO and European Union member states, facilitating the identification of differences in maturity and areas where interoperability needs to be strengthened.

A preliminary application of the model to the case of Romania suggests that the existing infrastructure falls between levels II and III. Current systems allow for multi-sensor integration and possess significant surveillance capabilities; however, the extensive use of artificial intelligence for predictive analysis, the automatic adaptation of defence systems and the integration of algorithmic governance are still at an early stage. Consequently, the main direction of development should not be the proliferation of sensors, but rather the evolution towards an integrated cognitive architecture capable of transforming data into rapid, explainable decisions that comply with the requirements of international law.

Looking ahead, MARMM has the potential to become a benchmark tool for assessing the digital progress of maritime organisations, providing a common analytical methodology in a field where technological transformation is advancing faster than the development of doctrinal concepts and institutional mechanisms. In this regard, the model does not aim to classify technologies, but rather to measure the maturity of the decision-making architectures that underpin 21st-century maritime security.

8.2. Strategic recommendations for Romania on the development of an artificial intelligence-assisted maritime architecture

The analysis carried out in the preceding chapters highlights the fact that Romania possesses a significant institutional and technological foundation for the development of a modern maritime security architecture. The existence of the SCOMAR system, membership of NATO and the European Union, the strategic ‘ ’ role of the Port of Constanța, the development of offshore energy infrastructure and the experience gained in monitoring the maritime domain constitute significant advantages. However, the rapid transformation of the security environment in the Black Sea basin necessitates a shift from an architecture predominantly focused on surveillance to one based on anticipation, information integration and decision-making support assisted by artificial intelligence.

This transformation does not entail abandoning existing infrastructure, nor does it involve replacing current systems with entirely new platforms. Such an approach would be difficult to justify from both an economic and an operational perspective. The recommended course of action is that of progressive modernisation, achieved by integrating existing capabilities into a unified cognitive architecture capable of harnessing the benefits of emerging technologies without compromising operational continuity.

In this context, the development of a national strategy on the use of artificial intelligence in maritime security should be the primary institutional objective. At present, the various components of the maritime security system are developed according to the needs of each institution, without a common conceptual framework regarding the role of artificial intelligence in the decision-making process. Drawing up such a strategy would enable the definition of unified objectives regarding interoperability, data standardisation, algorithm governance and the gradual integration of autonomous platforms into the national security architecture.

At the same time, it is necessary to develop a national multisensory fusion infrastructure. Existing systems already generate very large volumes of information, but this is analysed within different institutional frameworks and uses technical standards that do not always allow for the automatic exchange of data. The creation of a national information integration platform, capable of correlating in real time data from radars, AIS systems, optoelectronic sensors, autonomous vehicles, satellite imagery and cyber sources, would represent one of the most important investments in enhancing maritime resilience.

A natural step in this direction is the implementation of pilot projects based on autonomous maritime vehicles. Romania should not aim for the immediate development of a large fleet of autonomous platforms, but rather for their operational validation under the specific conditions of the Black Sea. An experimental programme conducted in the area of the Port of Constanța and its offshore infrastructure would enable the performance of autonomous systems to be assessed under real-world conditions of traffic, radio propagation and the impact of electronic warfare. The results of these projects could subsequently form the basis for the gradual expansion of the proposed architecture.

A distinct priority must be given to developing national capabilities in the field of explainable artificial intelligence. At present, attention is focused almost exclusively on the performance of detection and classification algorithms. However, as these systems become involved in the decision-making process, transparency and the ability to audit algorithmic recommendations will become just as important as technical accuracy. Romania should promote the development of platforms capable of explaining to decision-makers the reasons why a particular contact is classified as a threat and of maintaining a verifiable record of the entire decision-making process.

The transformation of maritime architecture must also be accompanied by investment in human resources. The introduction of artificial intelligence does not reduce the need for specialised personnel, but rather changes the profile of the skills required. Operators of maritime systems will need to understand both the functioning of sensors and naval platforms, as well as the principles of algorithmic analysis, probabilistic assessment and digital governance. Consequently, professional training programmes for naval units, the Coastguard and other institutions involved in maritime

security should include modules dedicated to artificial intelligence, data analysis, cyber security and multi-domain interoperability.

A further strategic direction is the development **of** a **National Centre for Artificial Intelligence and Maritime Security**, organised as a joint platform for research, testing and operational validation. Such a centre could bring together institutions with responsibilities in the field of maritime security, universities, research institutes and industrial partners, facilitating the transfer of research results to operational applications. At the same time, it could become a point of reference for cooperation with NATO and European Union bodies involved in the development of emerging maritime technologies.

From the perspective of international cooperation, Romania should aim for active integration into European and Euro-Atlantic programmes concerning the protection of critical maritime infrastructure. Participation in joint projects on the use of autonomous platforms, the exchange of maritime data, the development of standards for artificial intelligence and interoperability testing would accelerate the modernisation of national capabilities and reduce the costs associated with the independent development of complex technological solutions. In this regard, Romania's advantage lies not only in its geographical position on the Black Sea, but also in its potential to become an operational testing ground for architectures tailored to the security environment on NATO's eastern flank.

At the same time, strengthening maritime security must be accompanied by the development of an appropriate legal framework for the use of artificial intelligence. The introduction of autonomous systems and decision-support algorithms raises issues concerning liability, auditability, data protection, cyber security and the use of force. The development of clear rules on the certification of algorithms used in the field of maritime security, standards of explainability and obligations regarding the preservation of digital evidence would help to increase institutional trust and strengthen the legitimacy of the use of these technologies.

In the medium term, Romania could consider developing a **'digital twin'** of the Port of Constanța and the adjacent maritime infrastructure. Such a system would enable the real-time simulation of various risk scenarios, the testing of operational responses and the training of artificial intelligence algorithms without affecting the actual infrastructure. The integration of data from sensors, autonomous platforms and predictive models would transform the digital twin into an essential tool for operational planning and the assessment of the resilience of critical infrastructure.

From a long-term perspective, Romania's strategic objective should be to progress towards a level of maturity corresponding **to Level V of the Maritime AI Readiness Maturity Model (MARMM)**. This involves not only the introduction of new technologies, but the transformation of the entire security architecture into an adaptive cognitive system, characterised by the integration of artificial intelligence, resilient communications, autonomous platforms and a decision-making process centred on human responsibility, in accordance with the **Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)** model.

In conclusion, the modernisation of Romania's maritime security should not be viewed as an exclusively technological project. It represents a complex process of institutional, doctrinal and legal transformation, in which artificial intelligence acts as a multiplier of existing capabilities rather than a substitute for them. Romania's strategic advantage will not depend on the number of autonomous systems acquired or the complexity of the algorithms implemented, but on its ability to integrate these technologies into a coherent, interoperable architecture that is compatible with the requirements of international law and Euro-Atlantic collective security.

CHAPTER 9

Limitations of the research and future directions

Any research analysing the impact of emerging technologies on maritime security is influenced by the rapid pace of technological evolution and the dynamic nature of the international strategic environment. Consequently, the conclusions of this study must be interpreted in the light of these particularities and the inherent limitations of an approach that seeks both to analyse current realities and to formulate forward-looking conceptual models.

A first limitation stems from the forward-looking nature of a significant part of the analysis. Although the research draws on examples and lessons from recent conflicts, particularly in the Black Sea region and from attacks on critical maritime infrastructure, the proposed conceptual models – **the Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)**, **the Adaptive Maritime Security Zone (AMSZ), the Dynamic Threat Index (DTI)** and **the Maritime AI Readiness Assessment Model (MARMM)** – are theoretical constructs designed to organise and optimise the decision-making process. These do not constitute architectures that have been fully implemented and validated within an existing operational system and, therefore, require further verification through experimental exercises, simulations and real-world applications.

A second limitation concerns access to operational information. Many of the technologies used in the field of maritime security, particularly those developed for military and intelligence organisations, are classified or only partially accessible in the public domain. For this reason, the analysis was based on information drawn from official documents, published doctrines, institutional reports, scientific literature and lessons learnt from recent conflicts. It is possible that certain capabilities already under development or in operation may not be fully reflected in the sources accessible to the research.

Another limitation stems from the evolving nature of artificial intelligence. The performance of algorithms depends directly on the quality of the data used for training, the continuous updating of models, and the operational context in which they are deployed. An algorithm that performs well in a controlled environment may produce different results in a theatre of operations characterised by electromagnetic jamming, degraded communications, severe weather conditions or deliberate attempts to mislead. Consequently, no conceptual model can completely eliminate the uncertainty inherent in the decision-making process, and artificial intelligence must be viewed as a risk-mitigation tool rather than a foolproof mechanism.

The study is also limited by the fact that it does not include experimental validation of the proposed algorithms. The concepts relating to dynamic threat assessment, automatic adaptation of security zones or the integration of cognitive architectures were developed on the basis of established principles of artificial intelligence and multisensory fusion, but have not been tested in a fully functional operational environment. A subsequent stage of research should aim to implement these models in dedicated simulators or in exercises conducted in conjunction with naval units and the relevant maritime authorities.

A further limitation concerns the legal dimension of the research. Although this study proposes the concept of **'cognitive due diligence'** and analyses the implications of the use of artificial intelligence for states' obligations regarding the protection of critical maritime infrastructure, international law does not currently contain explicit rules governing in detail the use of artificial intelligence in the field of maritime security. The evolution of state practice, the development of international case law and the consolidation of technical standards will inevitably influence the interpretation of existing obligations and may lead to the emergence of new requirements regarding the responsible use of these technologies.

Furthermore, this research does not aim to provide a comparative assessment of all maritime security architectures developed globally. The analysis focuses on the Euro-Atlantic context and, in particular, on the implications for Romania, NATO and the European Union. Consequently, the specific characteristics of regions such as the Indo-Pacific, the Persian Gulf or the South China Sea – where the strategic environment and institutional architectures differ significantly – have not been analysed in detail and may form the subject of separate research.

These limitations do not diminish the relevance of the conclusions drawn, but rather define the framework within which they must be interpreted. The main aim of the study was not to develop

a technological product or to experimentally validate a command and control system, but to formulate a conceptual framework capable of integrating technological, operational and legal developments into a coherent vision of the future of maritime security.

From this perspective, the research opens up several avenues for further investigation. A first avenue concerns the development and experimental validation of the concept **of an Adaptive Human- e Maritime Decision-Making Architecture (AHMDA)** through the use of simulation platforms and multi-domain exercises. A second direction involves the development of mathematical models to quantify **the Dynamic Threat Index (DTI)** and their integration into predictive analysis systems based on machine learning.

A particularly promising area of research is the development of **digital twins** for ports and critical maritime infrastructure, capable of replicating the operation of physical systems in real time and enabling the simulation of complex attack scenarios, the testing of operational responses and the optimisation of decision-making processes without affecting the actual infrastructure.

At the same time, future research should analyse the legal implications of the use of explainable artificial intelligence (*XAI*), the auditability of algorithms and the allocation of responsibility within AI-assisted decision-making architectures. The evolution of these technologies will inevitably require the adaptation of standards relating to international liability, the duty of care, cooperation between states and the protection of critical infrastructure.

Last but not least, the **Maritime AI Readiness Maturity Model (MARMM)** can serve as the basis for comparative research into the maturity levels of various national maritime architectures. Its application within NATO member states, the European Union or other maritime regions would enable the development of objective indicators regarding the degree of integration of artificial intelligence and could help underpin public policies and investment strategies adapted to technological developments.

Overall, the limitations identified confirm that maritime security supported by artificial intelligence is a field undergoing profound transformation, in which technological development, the adaptation of the legal framework and the evolution of operational doctrines must be analysed together. It is precisely this interdependence that justifies the need for an interdisciplinary approach and opens up important avenues for future research into cognitive architectures for maritime security.

9.1. Original contributions of the study

Contemporary literature on maritime security is paying increasing attention to the use of artificial intelligence, autonomous platforms and the digitalisation of surveillance processes. However, most existing research analyses these developments from a predominantly technological perspective, focusing on sensor performance, the development of autonomous vehicles or the integration of IT systems into naval architectures. By contrast, the legal implications of these transformations, the organisation of decision-making processes, and the relationship between artificial intelligence, human responsibility and states' international obligations are analysed to a considerably lesser extent.

Building on this observation, this study proposes an interdisciplinary approach that integrates perspectives from international law, strategic studies, maritime security and artificial intelligence within a unified conceptual framework. The aim of the research is not merely to describe emerging technologies, but to develop a conceptual architecture capable of organising the decision-making process in the field of critical maritime infrastructure protection and of providing tools for assessing the maturity of artificial intelligence-based systems.

The study's first original contribution lies in the formulation **of the** concept **of the Human-Centred Adaptive Maritime Decision-Making Architecture (AHMDA)**. This represents a conceptual command and control model in which artificial intelligence, autonomous systems and multisensory fusion are integrated into an architecture centred on human responsibility. Unlike approaches that

seek to expand the autonomy of algorithms, AHMDA is based on the premise that the use of artificial intelligence must enhance human decision-making rather than replace it. The model proposes a redistribution of cognitive functions between the operator and the algorithm, whilst maintaining human control over all decisions that have significant legal or operational consequences.

Another contribution is the development of the concept of **the Adaptive Maritime Security Zone (AMSZ)**. In the specialist literature, maritime security zones are usually defined by fixed geographical boundaries and standardised protection measures. This study proposes a reinterpretation of this concept by transforming the security zone into a dynamic space, capable of constantly adapting its configuration in line with predictive risk assessments and changes in the operational environment. The AMSZ thus introduces an adaptive dimension that allows for the redeployment of resources and the reorganisation of the defensive posture before a threat materialises.

The third contribution is the development **of the Dynamic Threat Index (DTI)**, a model for the continuous and probabilistic assessment of maritime threats. Unlike traditional binary classifications, the DTI integrates information from multiple sources and allows for the continuous updating of the risk level associated with each detected contact. This model aims to support the decision-making process by prioritising threats and reducing the time required to adopt operational measures.

From a legal perspective, the study proposes the concept of **Cognitive Due Diligence**, which extends the analysis of the international duty of due diligence in the context of the use of artificial intelligence. The concept is based on the idea that technological developments influence the content of the obligations of prevention and protection undertaken by states. Without creating a new, autonomous legal obligation, Cognitive Due Diligence describes the emerging cognitive dimension of the duty of care, expressed through the reasonable and proportionate use of intelligent tools to identify and anticipate threats to critical maritime infrastructure.

Another original contribution is the development **of the Maritime AI Maturity Model (MARMM)**, a conceptual model designed to assess the level of maturity of artificial intelligence implementation within maritime security architectures. MARMM goes beyond assessments based solely on the degree of digitalisation and proposes an integrated analysis that includes technology, interoperability, decision-making, legal governance and organisational resilience. The model provides a useful tool both for the comparative assessment of different national architectures and for underpinning modernisation strategies.

A further contribution of this research lies in the integration of all these concepts into a unified analytical model. Instead of a fragmented approach, in which artificial intelligence, maritime infrastructure, decision-making and legal obligations are analysed separately, this study proposes a systemic perspective on maritime security. In this view, technology is not an end in itself, but an element of a cognitive architecture in which information, analysis, decision-making and legal responsibility function as parts of the same process.

The originality of the research is further reinforced by the development of an operational scenario applied to the Port of Constanța, used for the conceptual validation of the proposed architecture. The scenario does not aim to replicate a real incident, but demonstrates how artificial intelligence, autonomous platforms, multisensory fusion and human control can operate in an integrated manner within an operational environment characterised by multi-domain threats and short reaction times.

Overall, the contribution of this study goes beyond the analysis of emerging technologies and proposes a conceptual framework for the organisation of future maritime security architectures. By integrating technological, operational and legal dimensions, the research aims to provide both an analytical tool for the academic community and a potential benchmark for the development of doctrines and public policies concerning the protection of critical maritime infrastructure.

Table 1. Original contributions of the study

Contribution	Type	Purpose
AHMDA	Conceptual model	Organisation of AI-assisted maritime decision-making
AMSZ	Operational concept	Dynamic adaptation of maritime security zones
DTI	Analytical tool	Probabilistic and continuous threat assessment
Cognitive Due Diligence	Legal concept	Reinterpreting the duty of care in the context of AI
MARMM	Maturity model	Assessment of the level of AI integration in maritime security
The AHMDA Integrated Model	Interdisciplinary framework	Integration of technological, operational and legal dimensions

9.2. Methodological basis for the concepts developed in the study

1. General considerations

The concepts developed in this study have been formulated as conceptual models designed to analyse and organise the decision-making process regarding the protection of critical maritime infrastructure in a context characterised by the integration of artificial intelligence and autonomous systems.

These models do not constitute technical standards, official doctrines of NATO or the European Union, nor do they represent legal categories established in international law. They are analytical tools proposed to facilitate an understanding of the transformations brought about by artificial intelligence in contemporary maritime security architectures.

The methodology used to develop these concepts was based on integrating legal research, strategic studies, the analysis of complex systems and the literature on artificial intelligence, with the aim of identifying functional relationships that have been insufficiently explored in the specialist literature.

2. Stages of concept development

The process of developing the conceptual models involved five successive stages.

In the first stage, a critical analysis was carried out of the specialist literature on maritime security, the protection of critical infrastructure, artificial intelligence, autonomous systems and international law. Academic works, military doctrines and documents produced by NATO, the European Union, the International Maritime Organisation (IMO), the United Nations and other relevant institutions were analysed.

In the second stage, the main conceptual gaps in the current literature were identified. The analysis highlighted that most research addresses the technological, operational or legal dimensions separately, without developing an integrated model of the decision-making process assisted by artificial intelligence.

The third stage involved the development of individual conceptual models capable of addressing these gaps. Each concept was constructed on the basis of established theories, but by integrating and extending them into a new architecture applied to maritime security.

In the fourth stage, the functional relationships between the concepts were analysed, with the aim of avoiding overlaps and ensuring their complementarity.

The final stage involved integrating all the models into a unified doctrinal framework, defined in the study as the concept of maritime cognitive superiority.

3. The genesis of each concept

3.1 Adaptive, human-centred maritime decision-making architecture (AHMDA)

The AHMDA concept was developed by integrating several existing doctrinal approaches:

- Human-Centred Artificial Intelligence;
- Human-in-the-Loop;
- Explainable Artificial Intelligence;
- Decision-Centric Warfare;
- Joint All-Domain Command and Control;
- Maritime Command and Control.

The originality of the model lies in the integration of these approaches into a single architecture designed for the decision-making process regarding the protection of critical maritime infrastructure, and in the explicit inclusion of legal liability as a structural element of the system.

3.2 Adaptive Maritime Security Zone (AMSZ)

The AMSZ is based on classical theories concerning:

- Maritime Security Zones;
- Port Security;
- Maritime Domain Awareness;
- Dynamic Risk Assessment.

The study's contribution lies in transforming the security zone from a static geographical perimeter into an adaptive operational structure, continuously reconfigured on the basis of predictive assessments carried out using artificial intelligence.

3.3 Dynamic Threat Index (DTI)

The DTI was developed based on:

- classical risk assessment models;
- probabilistic risk assessment;
- Bayesian reasoning;
- machine learning;
- multisensor fusion.

The originality lies in the continuous integration of all information sources into a single index designed to support maritime decision-making.

3.4 The Maritime AI Readiness Assessment Model (MARMM)

MARMM is based on:

- the Capability Maturity Model;
- the AI Readiness Index;
- Digital Maturity Models.

The proposed model adapts these methodologies to the field of maritime security and introduces additional criteria relating to interoperability, governance, explainability and human control.

3.5 Cognitive Due Diligence

This concept derives from:

- the international duty of due diligence;
- the principle of prevention;
- the positive obligations of states;
- standards on the responsible use of artificial intelligence.

The contribution of this research lies in expanding the cognitive dimension of the duty of due diligence, arguing that technological development influences the content of the duty of prevention without altering its legal nature.

3.6 Maritime Cognitive Resilience

The concept draws on the literature concerning:

- cyber resilience;
- operational resilience;
- cognitive resilience;
- resilience engineering.

The originality lies in defining the resilience of the cognitive process of the entire maritime architecture, rather than just the IT infrastructure.

4. Integration of models

The concepts developed do not function independently.
They form a unified conceptual architecture.
AHMDA organises the decision-making process.
AMSZ organises the operational space.
DTI provides a probabilistic assessment of threats.
MARMM measures the degree of institutional maturity.
Cognitive Due Diligence underpins legal obligations.
Maritime Cognitive Resilience safeguards the functioning of the entire system.
Together, these form the concept of maritime cognitive superiority.

5. The nature of the concepts

The authors emphasise that the models developed in this study represent **original conceptual constructs**, devised with the aim of organising and explaining a field undergoing rapid transformation.

These concepts are not presented as official doctrines, nor as existing normative standards. They constitute theoretical tools intended to facilitate future research, doctrinal development and the formulation of public policies regarding the responsible use of artificial intelligence in the field of maritime security.

The originality of the research does not stem from the isolated use of terms such as *'human-centred AI*', *'due diligence*', *'AI readiness'* or *'cognitive resilience*' – which are already well-established in the specialist literature – but from their integration into a coherent conceptual framework and the development of new models tailored to the protection of critical maritime infrastructure.

CHAPTER 10

The vulnerabilities of artificial intelligence in maritime security and the need for cognitive resilience

The integration of artificial intelligence into modern maritime security architectures offers undeniable advantages in terms of the speed of information processing, early threat detection and the optimisation of decision-making. However, these benefits are accompanied by the emergence of new vulnerabilities that did not exist in traditional command and control systems. To the extent that operational decision-making relies increasingly on algorithmic analysis, system security is no longer determined solely by the protection of physical infrastructure or IT networks, but also by the integrity of the artificial intelligence models that process the information.

This shift alters the very nature of operational risk. In the past, compromising a radar or a command centre generally involved the physical destruction of equipment or the disruption of communications. In AI-assisted architectures, a system may continue to function apparently normally, yet generate erroneous conclusions as a result of data manipulation, algorithm degradation or the tampering of the machine learning process. From this perspective, the

vulnerability is no longer always visible, and its effects may only become apparent when a wrong decision is made.

One of the most significant challenges is posed by **adversarial** *attacks* **on artificial intelligence models**. These aim to subtly modify the input data so that the algorithm misclassifies an object or behaviour. In the maritime environment, such attacks can lead to a naval drone being mistakenly identified as a civilian vessel or, conversely, a legitimate contact being classified as a threat. The insidious nature of these techniques lies in the fact that the alterations introduced may be almost imperceptible to the human operator, yet sufficient to influence the outcome of algorithmic processing.

A distinct vulnerability is **the tampering with the datasets used to train the algorithms** (*data poisoning*). If an adversary succeeds in introducing erroneous data into the machine learning process, the system's performance may be degraded in the long term, without apparently affecting its operation. In the field of maritime security, such an attack could gradually alter the classification criteria for certain types of contacts, reducing the system's ability to identify real threats or significantly increasing the number of false alarms.

Added to these risks are the vulnerabilities arising from **the manipulation** of **data sources**. Modern maritime security systems simultaneously utilise information from radars, AIS systems, satellite imagery, acoustic sensors and autonomous platforms. The compromise of one or more sources can have repercussions throughout the entire decision-making chain. For example, the falsification of AIS signals – a phenomenon already documented in numerous maritime regions – can lead to the creation of an erroneous operational picture if the algorithms are unable to identify inconsistencies between the various sources of information. Similarly, **spoofing** techniques applied to global navigation satellite systems can alter the apparent position of an autonomous platform, affecting both the navigation process and threat assessment.

Electronic warfare presents a further challenge. In contemporary conflicts, jamming of communications and disruption of the electromagnetic spectrum are frequently used to degrade surveillance and command systems. In the case of architectures supported by artificial intelligence, the effects of these actions are amplified, as a reduction in the quality of input data can simultaneously affect the performance of all algorithms that utilise that information. Thus, the resilience of artificial intelligence depends directly on the resilience of the information infrastructure that supports it.

Nor should the inherent limitations of machine learning algorithms be overlooked. Even in the absence of deliberate attacks, artificial intelligence models may exhibit **algorithmic bias**, resulting from an unbalanced distribution of the data used for training or from implicit assumptions embedded in the development process. In the field of maritime security, this phenomenon can lead to the overestimation or underestimation of certain categories of threats, affecting the balance of the decision-making process. For this reason, the performance of algorithms must be continuously assessed, and models must be recalibrated in line with changes in the operational environment.

Another vulnerability is linked **to operators' excessive reliance on algorithmic recommendations**, a phenomenon known in the specialist literature as *'automation bias*'. As intelligent systems demonstrate a high level of accuracy, there is a risk that operators will automatically accept the conclusions generated by algorithms without conducting an independent critical analysis. In crisis situations, this tendency can lead to a reduction in effective human control over the decision-making process and to the operator becoming merely an executor of the system's recommendations.

The opposite phenomenon, known as *'algorithm aversion'*, can have equally negative effects. A lack of trust in algorithmic recommendations may lead to correct warnings being ignored and delays in taking the necessary measures. Consequently, the effectiveness of a maritime architecture assisted by artificial intelligence does not depend solely on the technical performance of the algorithms, but also on the development of a balanced relationship between human expertise and automated analysis.

These vulnerabilities demonstrate that the introduction of artificial intelligence does not eliminate uncertainty, but rather transforms it. Whilst traditional systems were primarily exposed to physical and cyber threats, cognitive architectures are also vulnerable to the manipulation of the learning process, information flows and classification mechanisms. Consequently, the security of a maritime infrastructure assisted by artificial intelligence cannot be assessed solely on the basis of algorithmic performance, but must be analysed in terms of the entire architecture's capacity to identify, absorb and correct these disruptions.

In this context, this study proposes the introduction of the concept of **Maritime Cognitive Resilience (MCR)**. This can be defined as **the capacity of an artificial intelligence-assisted maritime architecture to maintain its essential functions of observation, analysis and decision support even in the face of deliberate or accidental degradation of the data, algorithms or information infrastructures that underpin the cognitive process**.

Unlike cyber resilience, which is predominantly focused on the protection of networks and IT systems, cognitive resilience concerns the protection of the decision-making process as a whole. It entails the existence of multi-sensory validation mechanisms, continuous verification of information consistency, automatic anomaly detection, the use of explainable algorithms, the ability to audit decisions, and the maintenance of human control over critical stages of the operational process.

From this perspective, the concept **of Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)** takes on an additional function. It is no longer merely a model for integrating artificial intelligence into the decision-making process, but also a mechanism for strengthening cognitive resilience. The distribution of functions between algorithms and human operators, the use of multisensory convergence and the integration of principles of explainability reduce the likelihood that an isolated vulnerability will compromise the entire decision-making process.

In conclusion, the development of AI-assisted maritime architectures must be accompanied by the development of mechanisms to protect artificial intelligence itself. Cognitive superiority cannot exist without cognitive resilience, and the performance of algorithms cannot replace the need for robust governance, human control and continuous verification of the decision-making process. In the future, the strategic competitiveness of states will depend not only on the ability to develop more efficient algorithms, but also on the ability to build cognitive systems capable of withstanding manipulation, uncertainty and multi-domain attacks.

10.1. Maritime cognitive superiority – a new vision of maritime security in the 21st century

The evolution of digital technologies, the proliferation of autonomous systems and the integration of artificial intelligence into command and control architectures are bringing about a profound change in the way maritime superiority is understood. For centuries, naval power has been assessed using predominantly material indicators: fleet size, ship tonnage, firepower, control of sea lanes, or the ability to project power over long distances. These elements remain relevant today, but are no longer sufficient to explain strategic advantage in an environment characterised by automation, interconnectivity and the unprecedented acceleration of the information cycle.

Recent conflicts demonstrate that many maritime operations are decisively influenced not by material superiority per se, but by the ability of the actors involved to identify, interpret and utilise information before the adversary does. In many situations, the success of an operation does not depend on the physical destruction of hostile platforms, but on the early detection of their intentions, the anticipation of their behaviour and the adoption of an appropriate decision within a shorter timeframe than that available to the adversary. In this context, strategic advantage is increasingly determined by what can be defined as **cognitive superiority**.

In this study, the concept of **maritime cognitive superiority** is used to describe the ability of a state or organisation to transform information derived from the maritime environment into operational knowledge, legal assessment and strategic decision-making at a faster pace than that of

the adversary. This definition shifts the emphasis from the mere accumulation of information to the quality of the process by which it is analysed, integrated and utilised in decision-making.

Cognitive superiority must not be confused with informational superiority. A state may possess a very large number of sensors, satellites and autonomous platforms, yet still be unable to make effective decisions rapidly if the information collected is not integrated into a coherent architecture. Similarly, a small volume of information, if processed efficiently and contextualised appropriately, can generate a decisive strategic advantage. Cognitive superiority is, therefore, the result of transforming information into knowledge and knowledge into legitimate action.

This transformation requires the existence of an information ecosystem capable of functioning as a unified organism. Sensors, autonomous platforms, communications, artificial intelligence algorithms and human operators can no longer be analysed as independent elements. The value of each component stems from its ability to help reduce uncertainty and accelerate the decision-making process. From this perspective, technology is not the ultimate goal of the transformation, but rather the cognitive infrastructure upon which operational superiority is built.

The concept of **Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)** proposed in this study represents the organisational expression of this cognitive superiority. Whilst cognitive superiority defines the strategic objective, AHMDA describes the institutional mechanism through which it can be achieved. The architecture redistributes cognitive functions between algorithms and human operators, so that each decision benefits simultaneously from the speed of automated analysis and the legal and strategic judgement of the human element.

Complementarily, **the Adaptive Maritime Security Zone (AMSZ)** provides the spatial dimension of cognitive superiority. The security zone is no longer a static geographical boundary, but a dynamic space that continuously adapts in line with changes in the probability of threats. The redeployment of sensors and autonomous platforms is no longer determined exclusively by events that have already occurred, but also by the anticipation of changes in the operational environment.

In turn, **the Dynamic Threat Index (DTI)** provides the analytical tool through which disparate information is transformed into a probabilistic risk assessment. Instead of rigid, binary classifications, the DTI allows for a continuous assessment of the threat level and the ongoing adaptation of the operational response. This mechanism reduces uncertainty without removing the need for human oversight of the final decision.

However, cognitive superiority cannot exist in the absence of adequate governance. For this reason, the concept of **Cognitive Due Diligence** complements the legal dimension of the proposed architecture, demonstrating that technological development influences the standard of due diligence required of states. The use of artificial intelligence does not alter the fundamental nature of international obligations, but it does change the level of rigour required in organising prevention, risk assessment and the protection of critical maritime infrastructure.

Equally, **Maritime Cognitive Resilience (MCR)** ensures the continuity of the cognitive ecosystem's functioning. Superiority is defined not only by the speed of information processing, but also by the system's ability to withstand data manipulation, cyber-attacks, electronic warfare and attempts to compromise algorithms. A system that rapidly processes erroneous information does not generate superiority, but vulnerability. From this perspective, cognitive resilience becomes the indispensable condition for cognitive superiority.

Taken together, these concepts represent a paradigm shift in maritime security theory. The focus is shifting from the accumulation of platforms and sensors towards the development of an institutional capacity to integrate information, anticipate risks and make decisions that comply with both operational requirements and the demands of international law.

This concept has implications that extend beyond the technological sphere. It influences the way in which states organise their institutions, train their personnel, develop doctrines, invest in research and manage relations between civil and military authorities. Cognitive superiority thus becomes a characteristic of the entire security system and not merely of the technical infrastructure.

From this perspective, it can be argued that the 21st century marks the transition from the concept **of platform-based maritime power** to **that of knowledge-based maritime power**. Ships,

sensors and autonomous systems remain indispensable components of maritime security, but strategic advantage will no longer be determined solely by their individual performance. It will depend on states' ability to rapidly transform information into operational knowledge, knowledge into decision-making, and decisions into legitimate, proportionate and responsible action.

Consequently, the maritime superiority of the future must be understood first and foremost as **cognitive superiority**, and maritime architectures assisted by artificial intelligence must be designed not merely to observe more, but to understand better, decide more quickly and act in accordance with the principles of international law. It is precisely this conceptual shift that forms the theoretical foundation uniting all the models developed in this study and provides a direction for the evolution of future maritime security architectures.

CHAPTER 11

Towards a new doctrine of maritime security: integrating artificial intelligence, cognitive resilience and international law

The developments analysed in this study demonstrate that the transformations brought about by artificial intelligence go beyond the realm of technological modernisation and influence the conceptual foundations of maritime security. Whilst in recent decades the focus has been on the development of naval platforms, the expansion of surveillance systems and the enhancement of response capabilities, the emerging trend is shifting the centre of gravity towards the organisation of the decision-making process and the way in which information is transformed into operational knowledge.

This shift is not solely the result of technological progress. It reflects the changing nature of threats. Autonomous platforms, hybrid attacks, cyber operations and the use of artificial intelligence by state and non-state actors are reducing the time available for analysis and compressing the decision-making cycle to levels that can no longer be managed solely through traditional procedures. Under these circumstances, maritime architectures based on the progressive

accumulation of sensors and the centralisation of information in a single command centre are reaching their functional limits.

In this new reality, maritime security must be understood as an adaptive system, in which operational success depends on the relationship between technology, institutional organisation and legal norms. None of these components can, on its own, produce strategic superiority. Autonomous platforms without coordination mechanisms merely generate additional volumes of information. High-performance algorithms, when used in the absence of adequate governance, can accelerate the spread of errors. Similarly, legal norms, however well-developed they may be, cannot be effective unless they are integrated into the architecture of the decision-making process. Superiority stems exclusively from the coherent functioning of the entire system.

From this perspective, this study proposes a doctrine based on six complementary pillars. The first is **the Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA)**, which organises the distribution of cognitive functions between human operators and artificial intelligence. The second pillar is **the Adaptive Maritime Security Zone (AMSZ)**, through which the security space becomes an adaptive system rather than a rigid geographical demarcation. The third is **the Dynamic Threat Index (DTI)**, which transforms threat assessment into a continuous and probabilistic process. The fourth pillar is **the Maritime AI Readiness Assessment Model (MARMM)**, designed to assess organisational and technological progress. The fifth is **Cognitive Due Diligence**, which extends the analysis of the international duty of care in the context of artificial intelligence use. Finally, the sixth pillar is **Maritime Cognitive Resilience (MCR)**, which defines the system's ability to continue the decision-making process even in the face of deliberate degradation of information or digital infrastructure.

Together, these concepts form a doctrinal model in which artificial intelligence is no longer analysed as an isolated technology, but as part of an institutional and legal ecosystem. The originality of this model lies in the fact that it places the decision-making process at the centre of the security architecture, and technology is assessed in terms of its contribution to the quality and legitimacy of the decision, not solely on the basis of its technical performance.

This approach also has significant implications for how state responsibility is understood. In the traditional concept, responsibility was analysed primarily in relation to the actions taken after an incident had occurred. In the cognitive paradigm proposed in this study, the emphasis shifts towards the capacity of institutions to anticipate, organise and prevent risks. This shift brings maritime security closer to modern concepts of *anticipatory* governance, in which prevention becomes just as important as reaction.

From a doctrinal perspective, this development also necessitates a redefinition of the relationship between humans and technology. The literature on artificial intelligence is often dominated by the debate over the optimal level of autonomy for algorithms. This study proposes a different perspective. The fundamental issue is not how much autonomy artificial intelligence can be granted, but how collaboration between algorithms and the human element should be organised so that decision-making is simultaneously swift, efficient and legitimate. In this view, artificial intelligence is not a substitute for the commander, but a multiplier of their cognitive capacity.

This shift also has implications for institutional preparedness. Whilst the industrial concept prioritised the development of platforms and physical infrastructure, the cognitive paradigm emphasises the development of analytical skills, information interoperability and the strengthening of cooperation mechanisms between civil and military authorities. Investment in technology must be accompanied by investment in human capital, governance standards and mechanisms for auditing the decision-making process.

For Romania, adopting such a perspective is of particular strategic importance. Its geographical position on the Black Sea, the role of the Port of Constanța in the regional logistics architecture and the development of offshore energy infrastructure necessitate the establishment of a system capable of rapidly integrating information from the maritime, air, cyber and space domains. At the same time, membership of NATO and the European Union provides the necessary

institutional framework for the development of common standards on the responsible use of artificial intelligence and the strengthening of interoperability.

At an international level, the proposed model can contribute to the development of a common doctrine on the use of artificial intelligence in the protection of critical maritime infrastructure. Although each state is developing its own technological capabilities, the challenges posed by autonomous platforms, hybrid attacks and the vulnerability of subsea infrastructure are shared. In these circumstances, the development of common principles regarding algorithmic explainability, human oversight, auditability and information sharing may represent an essential element of collective security.

Overall, this study argues that the future of maritime security will not be determined solely by technological developments, but by states' ability to integrate technology into a coherent institutional and legal framework. Artificial intelligence, autonomous platforms and predictive analytics systems are not ends in themselves, but tools through which states can fulfil their obligations of protection, prevention and cooperation more effectively. It is precisely this integration of technology, law and institutional organisation that defines the new doctrine of maritime security proposed in this study.

CONCLUSIONS

The rapid transformations brought about by artificial intelligence, the development of autonomous systems and the expansion of critical maritime infrastructure demonstrate that maritime security is undergoing one of the most significant conceptual shifts of recent decades. Whereas, in the classical concept, the protection of maritime space was based on material superiority, geographical control of maritime routes and the performance of naval platforms, the analysis carried out in this study highlights the fact that strategic advantage is increasingly being determined by the ability to integrate information, anticipate threats and organise decision-making processes assisted by artificial intelligence.

The study demonstrates that the proliferation of unmanned naval vessels, the use of autonomous systems in maritime operations and the convergence of the maritime, cyber and space domains are fundamentally altering the architecture of maritime security. In this context, simply modernising sensors or increasing the number of surveillance platforms is no longer sufficient to ensure the effective protection of critical infrastructure. The main challenge no longer lies in gathering information, but in transforming it into a coherent operational picture and a decision taken quickly enough to respond to threats characterised by mobility, autonomy and low cost.

Building on this reality, the research proposes a paradigm shift in the understanding of contemporary maritime security. The main contribution of the study lies in the development of an integrated conceptual architecture that links the technological, operational and legal dimensions of the decision-making process. To this end, several conceptual models have been formulated and defined, designed to organise modern systems for the protection of critical maritime infrastructure.

The Adaptive Human-Centred Maritime Decision-Making Architecture (AHMDA) forms the core of this architecture, proposing a model in which artificial intelligence supports the decision-making process without removing human control over the use of force and the assumption of legal responsibility. The Adaptive Maritime Security Zone (AMSZ) redefines the security domain as an adaptive structure, configured according to the probabilistic evolution of threats rather than exclusively on the basis of fixed geographical boundaries. The Dynamic Threat Index (DTI) introduces a continuous risk assessment mechanism, whilst the Maritime AI Maturity Model (MARMM) provides a tool for assessing institutional maturity regarding the integration of artificial intelligence into the field of maritime security.

From a legal perspective, the study argues that the development of artificial intelligence also has implications for the international standard of due diligence applicable to states. The concept of Cognitive Due Diligence proposes a reinterpretation of the obligation to prevent and protect in light of new technological capabilities for anticipating and analysing risks, without altering the

fundamentals of international law concerning state responsibility. Complementarily, Maritime Cognitive Resilience highlights the need to protect not only physical and cyber infrastructure, but also the cognitive process through which information is transformed into decision-making.

The integration of these models leads to the formulation of a paradigm which the study defines as **maritime cognitive superiority**. From this perspective, strategic advantage no longer lies exclusively with the state that possesses the most ships, the most extensive network of sensors or the most advanced algorithms, but with the one that succeeds in transforming information into operational knowledge, knowledge into legitimate decision-making, and decision-making into effective action within a shorter timeframe than its adversary.

An analysis of recent incidents in the Black Sea and the vulnerabilities associated with critical maritime infrastructure confirms the practical relevance of this approach. For Romania, the development of an integrated architecture based on multisensory fusion, artificial intelligence and inter-institutional cooperation is a strategic necessity, given the role of the Port of Constanța in regional security, allied military mobility and the resilience of European logistics chains.

At the same time, the research highlights that technological development must be accompanied by the strengthening of the legal and institutional framework. The use of artificial intelligence in maritime security cannot be reduced to a technical issue, but entails requirements regarding algorithmic transparency, the auditability of decisions, human oversight and compliance with the principles of international law. The legitimacy of the decision-making process remains dependent on human accountability, even in an environment characterised by advanced automation.

The research findings also open up avenues for future developments. The proposed conceptual models can serve as a starting point for developing quantitative indicators to assess institutional maturity, for testing algorithms for dynamic threat assessment, and for integrating these into operational exercises conducted within NATO and the European Union. Equally, the concepts developed may contribute to the formulation of international standards on the responsible use of artificial intelligence in the protection of critical maritime infrastructure.

In conclusion, this study argues that 21st-century maritime security can no longer be explained solely through traditional concepts of naval power. The new generation of threats necessitates the development of cognitive architectures in which technology, human expertise and international law function as complementary elements of the same system. In this new paradigm, artificial intelligence is not the objective of the transformation, but the tool through which states can fulfil their obligations of protection, prevention and cooperation more effectively. The maritime superiority of the future will belong to those actors capable of transforming the speed of information into the legitimacy of decision-making, and the legitimacy of decision-making into sustainable security.

LIST OF ABBREVIATIONS

Abbreviation	Meaning
AHMDA	Adaptive Human-Centric Maritime Decision Architecture
AIS	Automatic Identification System
AI	Artificial Intelligence
AMSZ	Adaptive Maritime Security Zone
C2	Command and Control
CNN	Convolutional Neural Network
COP	Common Operational Picture
DTI	Dynamic Threat Index
EO	Electro-Optical
EU	European Union

Abbreviation	Meaning
GNSS	Global Navigation Satellite System
GRU	Gated Recurrent Unit
IR	Infrared
JADC2	Joint All-Domain Command and Control
LSTM	Long Short-Term Memory
MARMM	Maritime AI Readiness Maturity Model
MCR	Maritime Cognitive Resilience
MDO	Multi-Domain Operations
NATO	North Atlantic Treaty Organisation
OSINT	Open Source Intelligence
RCS	Radar Cross Section
SCMAR/SCOMAR	Maritime Traffic Control and Surveillance System (if the abbreviation is used in the study)
SCNR	Signal-to-Clutter-and-Noise Ratio
UAV	Unmanned Aerial Vehicle
USV	Unmanned Surface Vehicle
XAI	Explainable Artificial Intelligence

BIBLIOGRAPHY

I. Books

1. Aldrich R, *GCHQ: The Uncensored Story of Britain's Most Secret Intelligence Agency* (HarperPress 2010).
2. Booth K, *Navies and Foreign Policy* (Routledge 2014).
3. Gray CS, *The Leverage of Sea Power: The Strategic Advantage of Navies in War* (Free Press 1992).
4. Haykin S, *Neural Networks and Learning Machines* (3rd ed., Pearson 2009).
5. Heckman JJ and Rogers WH, *Artificial Intelligence and Decision Making* (Oxford University Press 2021).
6. Kaplan FD, *The Coming AI Revolution in Defence* (Yale University Press 2023).
7. Kraska J and Pedrozo RL, *International Maritime Security Law* (2nd ed., Brill Nijhoff 2022).
8. Russell S and Norvig P, *Artificial Intelligence: A Modern Approach* (4th ed., Pearson 2021).
9. Scharre P, **Army of None: Autonomous Weapons and the Future of War** (W W Norton 2018).
10. Till G, **Seapower: A Guide for the Twenty-First Century** (5th ed., Routledge 2024).

II. Articles

1. Aksenov V et al., 'Artificial Intelligence for Maritime Surveillance' (2024) *IEEE Access*.
2. Bellingham JG, 'Autonomous Maritime Systems' (2022) *Annual Review of Marine Science*.
3. Brundage M, 'Towards Trustworthy Artificial Intelligence' (2020) **Nature Machine Intelligence**.
4. Carlini N et al., 'Adversarial Machine Learning' (2021) *Communications of the ACM*.
5. Endsley MR, 'Situation Awareness in Dynamic Human Decision Making' (1995) *Human Factors*.
6. LeCun Y, Bengio Y and Hinton G, 'Deep Learning' (2015) *Nature*.

III. NATO documents

1. NATO, *Artificial Intelligence Strategy* (2021).
2. NATO, *Data Exploitation Framework Policy* (2023).
3. NATO Allied Maritime Command, *Maritime Strategy*.
4. NATO, *Emerging and Disruptive Technologies Strategy*.
5. NATO STO, *Artificial Intelligence for Military Decision Support*.
6. NATO CCDCOE, *Cyber Defence Handbook*.

IV. European Union documents

1. European Commission, *AI Act* (Regulation (EU) 2024/1689).
2. European Commission, *European Maritime Security Strategy*.
3. European Commission, *Action Plan on Military Mobility*.
4. European External Action Service, *EU Maritime Security Strategy* (2023).
5. European Union Agency for Cybersecurity (ENISA), *Threat Landscape* (2024).

V. IMO documents

1. International Maritime Organisation,
2. *SOLAS Convention*
3. *ISPS Code*
4. *MSC Guidelines on Maritime Autonomous Surface Ships (MASS)*
5. *Guidelines on Maritime Cyber Risk Management*

VI. UN Documents

1. United Nations Convention on the Law of the Sea (1982).
2. International Law Commission,
3. *Draft Articles on the Responsibility of States for Internationally Wrongful Acts* (2001).
4. UNESCO,
5. *Recommendation on the Ethics of Artificial Intelligence* (2021).
6. UN General Assembly,
7. *Global Digital Compact* (2024).

VII. Case law

1. *Pulp Mills on the River Uruguay (Argentina v Uruguay)* [2010] ICJ Rep 14.
2. *South Africa v Israel* (Provisional Measures, ICJ).
3. Advisory Opinion of the International Tribunal for the Law of the Sea on Climate Change (2024).
4. Advisory Opinion of the International Court of Justice on Climate Change (2025).

VIII. Technical reports

1. RAND Corporation,
2. *Artificial Intelligence and National Security*.
3. RAND,
4. *AI for Maritime Domain Awareness*.
5. MIT Lincoln Laboratory,
6. *Sensor Fusion for Maritime Surveillance*.
7. Centre for Naval Analyses (CNA),
8. *Autonomous Maritime Systems*.
9. DARPA,
10. *Sea Hunter Programme*.

IX. Standards and guidelines

1. ISO/IEC 23894:2023 — Artificial Intelligence — Risk Management.
2. ISO 31000:2018 — Risk Management.
3. NIST AI Risk Management Framework 1.0.
4. NIST Cybersecurity Framework 2.0.
5. OECD AI Principles.

X. Online resources

1. European Maritime Safety Agency (EMSA)
2. NATO Maritime Command (MARCOM)
3. NATO CCDCOE
4. International Maritime Organisation (IMO)
5. ENISA
6. European Commission
7. US Naval Institute
8. Royal United Services Institute (RUSI)
9. Chatham House
10. CSIS