MÁSTER UNIVERSITARIO EN INGENIERÍA INDUSTRIAL

TRABAJO FIN DE MÁSTER

Predictive Maritime Intelligence for Hybrid Threat Detection.

Autor: Gonzalo María de Araluce Bergel

Director: Alejandro González San Román

Co-Directora: Ileana Daniela Serban

Declaration of originality.

I declare under my responsibility that the Project presented with the title **"Predictive Maritime Intelligence for Hybrid Threat Detection"** at the ICAI School of Engineering of the Comillas Pontifical University in the academic year 2025/2026 is of my authorship and has not been presented previously for other purposes. The Project is not plagiarised from any other, either totally or partially, and the information that has been taken from other documents is duly referenced.

Use of Artificial Intelligence[1]

I declare under my responsibility that (indicate the correct option):

□ I have not used Artificial Intelligence in the preparation of this document.

■ I have used Artificial Intelligence in the preparation of this document and/or Annex B under the conditions allowed by Comillas Pontifical University, i.e. applying Level 2 of the Perkins et al. (2024) Assessment Scale: "AI can be used for pre-task activities such as brainstorming, description and initial research. This level focuses on the use of AI for planning, synthesising and generating ideas, but assessments should emphasise the ability to develop and refine these ideas independently". Specifically, Artificial Intelligence has been used to:

Claude code (Anthropic) and cursor were used as programming assistants, assisting with code generation and refinement. These applications aided in writing, optimizing and debugging Python scripts for the system described below. Additionally, AI was used for synthetic data generation for the machine learning classifiers as explained in the document.

AI models, specifically ChatGPT and google gemini were also used to find sources and research on relevant topics, as well as for proof reading and grammar revisions. These models were also used for the planning and scheduling of the project development.

[1] This declaration refers to the use of generative Artificial Intelligence to carry out the Project documents (Annex B and Memory). It does not apply to Projects where, by their nature, artificial intelligence must be used as part of them (application of machine learning techniques, neural networks, data analysis...).

Signature (student): 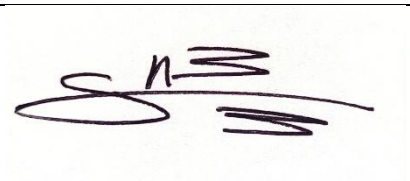
Date: 3rd July 2026.

Authorisation for Project delivery

Thesis supervisor	Thesis deputy supervisor (if any)
	Ileana Daniela Serban
Signature:	Signature: 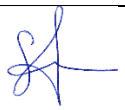
Date:	Date: 07/07/2026

MÁSTER UNIVERSITARIO EN INGENIERÍA INDUSTRIAL

TRABAJO FIN DE MÁSTER

Predictive Maritime Intelligence for Hybrid Threat Detection.

Autor: Gonzalo María de Araluce Bergel

Director: Alejandro González San Román

Co-Directora: Ileana Daniela Serban

Agradecimientos

Agradezco a mis directores, Alejandro González San Román y Ileana Daniela Serban, su tiempo y dedicación a este trabajo.

PREDICTIVE MARITIME INTELLIGENCE FOR HYBRID THREAT DETECTION.

Autor: De Araluce Bergel, Gonzalo María.

Directores:
Alejandro González San Román
Ileana Daniela Serban

Entidad Colaboradora: ICAI – Universidad Pontificia Comillas

ABSTRACT.

Hybrid threats in the maritime domain have become a central concern for international security, with grey zone operations exploiting legal complexity and traffic density of strategic chokepoints to work towards geopolitical goals while remaining below the threshold for armed conflict. This project sets to analyse the hybrid environment from a technical and geopolitical standpoint, while developing a predictive intelligence system, processing input data, derived from AIS among other sources through a data fusion pipeline based on the Joint Directors of Laboratories (JDL) model. Random Forest, Isolation Forest and XGBoost classifiers are used to detect anomalies and produce probabilistic threat assessments, as well as to infer behavioural intent. The system has been developed with a strong focus on explainability, for example using SHAP values to provide transparency at every stage, keeping a human-in-the-loop architecture in mind. This allows the developed system to be used as a toolkit for intelligence analysis rather than an automated analysis tool, supporting human judgement instead of replacing it. The case study has been focused on the Strait of Gibraltar as a critical chokepoint, while being extended to the space domain as a parallel case study to analyse its cross-domain applications and examine the future of hybrid operations. Both systems are evaluated against the project objectives, with the findings being used to formulate a set of policy recommendations addressed to decision makers at the European and NATO level.

Key words: Hybrid operations, maritime intelligence, machine learning, JDL data fusion, explainable AI, space domain awareness, grey-zone operations.

SUMMARY.

Hybrid maritime operations have evolved from a regional concern into a structural feature of the modern security environment. High profile incidents such as the sabotage of infrastructure such as the Nord Stream pipelines, the targeting of Baltic Sea data cables or the systematic use of dark vessel fleets across the Persian Gulf and North Atlantic have demonstrated that grey zone tactics are now an established instrument for statecraft. While the Strait of Gibraltar has not yet witnessed the same level of overt hybrid disruption as other areas of operation like the Baltic or South China Sea, its characteristics suggest that it is an inevitable future development zone for similar tactics. Firstly, it is traversed by more than 20% of global maritime trade and contains critical infrastructure connecting Europe to Africa and the Atlantic to the Mediterranean, making it of significant geopolitical importance. As a consequence, the region concentrates the conditions hybrid adversaries exploit, including dense legitimate traffic to mask illicit activity, overlapping jurisdictions and high value targets with limited security.

Current monitoring systems mostly identify isolated anomalies but rarely support intent inference or operational attribution, forcing a reactive approach to threats. This was the main capability gap identified, not as a function of insufficient data but of insufficient analytical architecture to contextualise threats. This project addresses that gap through the design, development and evaluation of a predictive maritime intelligence system built around data fusion criteria structured around the Joint Directors Laboratories framework. At JDL level 0, raw AIS messages are ingested via a live WebSocket connection and normalised alongside contextual data from sources such as vessel ownership registries, sanctions databases and port state control records. At level 1, this data is processed into a thirteen-dimensional feature vector capturing the vessel's kinematic behaviour, transmission anomalies or ownership data among other features for each tracked vessel. At level 2, three machine learning classifiers, a Random Forest, an Isolation Forest and

an XGBoost model are applied in parallel with their outputs being fused into a single probabilistic threat score. Additionally, a SHAP explainability layer computes the Shapley values for each assessment, ensuring that the contribution of each input feature is transparent and auditable by human analysts. Finally, a behavioural profile is generated for each vessel, assigning probability and confidence levels for behaviours such as shadowing or surveillance, while identifying the observations that support the inference in accordance with NATO STANAG 4559. A predictive geofencing model also runs in parallel to the behavioural inference system, forecasting potential geofencing breaches using Kalman filtering and Monte Carlo confidence intervals, with a network intelligence function identifying vessel ownership, data and potential links with other vessels in the area of operations, potentially identifying coordination between vessels.

The same framework is extended to the space domain as a parallel case study, demonstrating that the pipeline architecture can be generalised beyond maritime surveillance. Substituting input data for the domain equivalent, such as TLE as a replacement for AIS data, the space model tracks a selection of satellites highlighted by the Space Force and obtained via CelesTrak or SpaceTrack. Threat and confidence scores are similarly computed, applying the same fusion architecture and through the use of variables like anomalous manoeuvres being flagged.

The project concludes with an assessment against the original design objectives, policy recommendations to improve the state of readiness when facing hybrid operations and a critical account of the limitations and possibilities for future system development.

ABSTRACTO.

Las amenazas híbridas en el ámbito marítimo se han convertido en un aspecto central para la seguridad internacional, con operaciones en la "zona gris" que explotan la complejidad jurídica y el tráfico elevado en los puntos estratégicos de estrangulamiento para conseguir alcanzar objetivos geopolíticos manteniéndose por debajo del umbral del conflicto armado. Este proyecto analiza el entorno híbrido desde una perspectiva técnica y geopolítica, desarrollando al mismo tiempo un sistema de inteligencia predictiva que procesa datos de entrada derivados de fuentes de AIS entre otras, a través de una arquitectura basada en el modelo Joint Directors of Laboratories (JDL). Para detectar anomalías, se usan clasificadores tipo Random Forest, Isolation forest y XGBoost, produciendo evaluaciones de amenazas probabilísticas e inferiendo la intención de conducta. El sistema se ha desarrollado con un fuerte enfoque en la explicabilidad, utilizando por ejemplo valores SHAP para fomentar la transparencia, manteniendo una arquitectura de "humano en el bucle", o human-in-the-loop. Esto permite que el sistema desarrollado funcione como una herramienta de apoyo al análisis de inteligencia en lugar de pretender reemplazar el juicio humano. El estudio del caso se ha centrado en el Estrecho de Gibraltar como punto de estrangulamiento crítico, ampliándose al dominio espacial como caso de estudio paralelo para analizar sus aplicaciones en otros dominios y examinar el futuro de las operaciones híbridas. Por último, los sistemas se evalúan frente a los objetivos establecidos, utilizando los resultados para elaborar recomendaciones dirigidas a los responsables de la toma de decisiones tanto a nivel europeo como de la OTAN.

Palabras clave: Operaciones híbridas, inteligencia marítima, aprendizaje automático, fusión de datos (JDL), IA explicable, Conciencia del Dominio Espacial (SDA), operaciones en la zona gris.

RESUMEN.

Las operaciones marítimas híbridas han evolucionado de ser una preocupación regional a convertirse en una característica estructural del entorno de seguridad moderno.

Incidentes de alto perfil como el sabotaje de infraestructuras como los gasoductos Nord Stream, el ataque a cables de datos en el Mar Báltico o el uso sistemático de flotas de buques fantasma en el Golfo Pérsico o en el Atlántico Norte han demostrado que las operaciones híbridas ya son una herramienta consolidada en la política de Estado. Aunque el Estrecho de Gibraltar todavía no ha experimentado el mismo nivel de disrupción híbrida como otras áreas de operaciones como el mar Báltico o el mar de China, sus características sugieren que es una zona de desarrollo futuro inevitable para tácticas similares. En primer lugar, es transitado por más del 20% del comercio marítimo mundial y contiene infraestructuras críticas que conectan Europa con África y el Atlántico con el Mediterráneo, lo cual le otorga una gran importancia geopolítica. Como consecuencia, la región concentra las condiciones que suelen favorecer el desarrollo de operaciones híbridas, entre ellas, un trafico marítimo legítimo muy denso que facilita el enmascaramiento de actividad ilegal, jurisdicciones superpuestas y objetivos de alto valor con niveles de protección limitados.

Los sistemas actuales de vigilancia identifican principalmente anomalías aisladas, pero rara vez permiten inferir intenciones o la atribución operacional, lo que obliga a adoptar un enfoque reactivo ante las amenazas. Este era el principal vacío de capacidades identificado, proveniente de la falta de una arquitectura analítica insuficiente y no de la falta de datos suficientes. Este proyecto aborda esa brecha mediante el diseño, desarrollo y evaluación de un sistema de inteligencia marítima predictiva basado en la fusión de datos siguiendo el marco de referencia JDL. En el nivel 0 del modelo JDL, los mensajes AIS se ingieren a través de una conexión WebSocket en vivo y se normalizan junto con datos contextuales procedentes de fuentes como registros de propiedad de buques o bases de datos de sanciones. En el nivel 1, estos datos se procesan para formar un vector de características de trece dimensiones que refleja el comportamiento cinemático de cada buque monitorizado, así como otros datos como anomalías en transmisiones o información sobre su estructura de propiedad. En el nivel 2, se procesan los datos a través de tres clasificadores de aprendizaje automático, un Random Forest, un Isolation Forest y un modelo XGBoost, que se ejecutan en paralelo, fusionando sus resultados para asignar una probabilidad de amenaza. También se integra una función de explicabilidad basada en SHAP, que calcula los valores de Shapley para cada evaluación, garantizando que la contribución de cada input sea transparente y auditable por analistas humanos. Finalmente, se genera un perfil de comportamiento para cada buque, asignando niveles

de probabilidad y confianza para conductas como el shadowing o la vigilancia, resaltando al mismo tiempo las observaciones que sustentan los outputs de acuerdo con el estándar NATO STANAG 4559. A la vez, se ejecuta un modelo de geofencing predictivo que pronostica posibles intrusiones en perímetros usando el método de filtrado de Kalman e intervalos de confianza de Monte Carlo, junto a una función de inteligencia de redes que identifica la estructura de propiedad del buque, sus datos y los posibles vínculos con otros buques en el área de operaciones, identificando posibles indicios de coordinación entre ellos.

Este mismo marco se extiende a su vez al dominio espacial como un caso de estudio paralelo, demostrando que la arquitectura del sistema se puede generalizar más allá de la vigilancia marítima. Sustituyendo datos de entrada por sus equivalentes en el espacio (por ejemplo, usando TLE para reemplazar a los datos obtenidos del sistema AIS), el modelo rastrea una selección de satélites destacados por la Fuerza Espacial americana (U.S, Space Force) y obtenidos a través de CelesTrack o Space-Track. Las puntuaciones de amenaza y confianza se calculan de manera parecida, aplicando la misma arquitectura de fusión de datos y mediante el uso de variables como la señalización de maniobras anómalas.

El proyecto concluye con una evaluación frente a los objetivos de diseño originales, recomendaciones políticas para mejorar el estado de preparación al hacer frente a operaciones híbridas y un análisis crítico de las limitaciones y posibilidades para el desarrollo futuro del sistema.

Report index.

Chapter 1. INTRODUCTION ***14***

Chapter 2. STATE OF AFFAIRS. ***19***

Chapter 3. METHODOLOGY. ***102***

Chapter 4. SYSTEM CONCEPT. ***103***

Chapter 5. CONCLUSION. ***150***

Chapter 6. BIBLIOGRAPHY. ***157***

Chapter 1. INTRODUCTION

Hybrid threats in the form of maritime activity have become a growing security concern as grey-zone tactics such as infrastructure sabotage, covert intelligence collection and dual-use civilian vessels have spread from the South China Sea and the Baltic Sea to other strategic regions. These incidents demonstrate how adversaries can exploit complicated environments, both politically and geographically, to mask hostile behaviour under the cover of legitimate maritime activity. As these methods become more effective and widely adopted, they present an increasing risk to chokepoints like the Strait of Gibraltar, an area of strategic importance with dense commercial and military traffic, as well as a concentrated network of underwater infrastructure, creating ideal targets and conditions for concealment and operational deniability.

As incidents involving grey-zone operations, or activities designed to achieve political or military objectives without triggering armed conflict[2] have become more frequent, global concern over maritime hybrid threats has intensified. These operations often involve the sabotage of underwater critical infrastructure such as pipelines and cables, espionage

[2] NATO - Topic: Countering hybrid threats

disguised as commercial or research activity, or the use of civilian vessels such as flatbed tankers as platforms for other operations, including drone launches or electronic interference.

Some of the key advantages of these operations relate to the overall strategic efficiency, which allows actors to pursue their objectives while avoiding some of the downsides of conventional warfare such as the high cost, legal consequences and escalation risks[3]. As mentioned, one of their main advantages is their cost asymmetry, with damaging a subsea cable or disrupting commercial patterns being relatively inexpensive in comparison to the benefits obtained, creating an economic imbalance that benefits offensive behaviour[4]. The ambiguous nature of grey zone operations is also the key for actors to probe for reactions to their approaches, furthermore, in the context of alliances, ambiguity complicates establishing consensus decisions when member states have equal decisive power, potentially acting as a destabilising wedge. When behaviour remains plausibly deniable, political decision-makers remain cautious before responding and potentially escalating the situation. This can delay coordinated action and reduce the credibility of deterrence, creating an illusion where the political cost of response appears higher than the cost of action. This concept will be developed further as a method to quantify offensive and defensive offsets, setting a framework for the classification of potential objectives as a function of the operational efficiency of targeting infrastructure.

By improving the ability to detect and document suspicious maritime behaviour, AI-backed systems can have serious strategic implications. The main contribution, and the main objective explored with this project, is improving operational awareness to help with deterrence by removing the ambiguity hybrid operations need to function. By supporting signalling mechanisms and helping to build appropriate responses by collecting clear, non-deniable evidence, the aim is to contribute to raising the perceived risk of conducting hybrid operations, increasing deterrence outside of conventional warfare.

What makes the maritime domain particularly sensitive to the evolving dynamics is not only the relatively access to unguarded critical infrastructure, but also the structural

---

[3] Sheppard, L. R., & Conklin, M. (2019). *Warning for the Gray Zone*. https://www.csis.org/analysis/warning-gray-zone

[4] Coito, J. (2025). *Protecting Subsea Cables: Detect to Deter, Sue to Secure*. https://www.csis.org/analysis/protecting-subsea-cables-detect-deter-sue-secure

characteristics of the environment itself. Maritime spaces are saturated with legitimate activity and are legally complex, with activity taking place constantly within overlapping jurisdictions and regulatory frameworks. This complexity presents a cover for irregular behaviour to blend in, making it a challenge to distinguish between coincidence, negligence and intent. Furthermore, the growing development of an interconnected society has accentuated the criticality of chokepoints, which not only create a physical constraint to shipping routes, adding to the complexity of the navigational situation, but also serve as the natural path for critical infrastructure connecting states. Energy flows, digital connectivity and the basis of the goods and services economies that form the basis of the global economic system all rely on the security of chokepoints, increasing their vulnerability and the strategic leverage of targeting them. This dependency has been documented extensively, especially in the context of global energy and trade security, where disruptions to a small number of maritime corridors can produce disproportionate downstream economic effects[5][6]. The Strait of Gibraltar exemplifies this concentration of high-risk factors, connecting the Atlantic to the Mediterranean, hosting critical undersea cables and pipelines and with more than 20% of the world´s maritime trade crossing the shipping lanes in either direction[7], creating a high value target for maritime hybrid operations.

Another defining feature of grey-zone maritime activity is its incremental nature, as they build on seemingly minor deviations of standard operations to avoid detection and facilitate the approach to sensitive infrastructure. Individually, these actions might appear to be insignificant breaches of conduct but collectively, they might indicate adversarial interest or probing behaviour, laying the groundwork for sabotage. Therefore, the challenge is not only restricted to identifying anomalies, but understanding patterns over time and across multiple, reliable data sources. In this context, strengthening the analytical capacity to interpret maritime behaviour becomes relevant at a strategic level, beyond the operational and tactical aspects. Improving the ability to contextualise suspicious activity, document patterns and support evidence-based assessments can

[5] Komiss, W., & Huntzinger, L. (n.d.). *The Economic Implications of Disruptions to Maritime Oil Chokepoints*.

[6] Ulrichsen, K. C., & Krane, J. (n.d.). *Maritime Chokepoints and Risks to Global Shipping and Energy Security*.

[7] Mediterranean ports, hubs of global trade - We Build Value

reduce the level of ambiguity that serves as operational cover. This in turn, can allow for coordinated responses earlier in the escalation process, creating a more stable environment by supporting well grounded responses.

The project deals with an intersection of technological and strategic developments, responding to the growing hybrid threat and recognising that these are not isolated incidents, but part of a broader dynamic in which accurate and timely analysis can shape the outcome as much as "hard" deterrence. By focusing on how data-driven tools can support earlier detection and interpretation, the project seeks to analyse a vulnerability in the contemporary security sphere by applying developing technology to helping human perception, attribution and judgement. In order to achieve this, the system must satisfy a set of design objectives. First, it must draw data from sources that are realistically available rather than requiring privileged or classified access to ensure that the approach remains broadly applicable. Second, outputs must remain explainable at every stage, supporting rather than replacing human judgment, in line with the human in the loop philosophy explored further in chapter 4. Lastly, the system must be capable of ingesting live data to operate in near real time, since the operational value of the tool depends on supporting timely analyst decisions rather than retrospective review.

Once these objectives have been established, the methodology can be designed to ensure that the requirements are met. The initial phase of the project involves a review of the current state of affairs in maritime hybrid operations and the technological aspects enabling the development of such systems, identifying current capability gaps. Once these have been identified, the system concept is developed, addressing its philosophy, architecture and the data fusion and extraction pipeline that supports anomaly scoring models in chapter 4. Finally, the objectives met and limitations encountered are evaluated against the initial design, as well as including policy recommendations to address the issues encountered during the development of the system.

Additionally, one of the aims of this project is to evaluate the grey zone environment from an interdisciplinary standpoint. Not only is the framework designed to be applied across multiple areas of operation impacted by hybrid operations, such as space or the maritime domain, but also aims to develop an overarching view, including a non-technical perspective. Hybrid operations sit at the intersection between engineering, strategy and geopolitics, with conflict and legal ambiguity shaping grey zone activity, and strategic deterrence theory determining what kind of evidence and signalling actually changes

adversarial behaviour. The system is designed with this intersection in mind, rather than treating the legal and strategic dimensions as an afterthought to the technical work. The developed systems are available at "GonzaloMAB" for reference.

Chapter 2. STATE OF AFFAIRS.

2.1. State of affairs in hybrid operations.

Hybrid maritime operations have, until recently, predominantly constrained to the South China Sea[8] and the Baltic Sea[9], where hybrid fishing fleets, harassment incidents, and infrastructure sabotage have demonstrated how these tactics can be used to achieve strategic goals. In these regions, maritime activity has been used as a tool to apply political pressure while remaining below the threshold for conventional conflict through careful escalation management. For example, hybrid fishing fleets operating in contested waters have been used to reinforce territorial claims around disputed outposts such as Mischief or Whitsun Reef[10], while covert infrastructure sabotage and harassment incidents have demonstrated how maritime activity can be leveraged to influence the strategic environment. As these methods haven proven to be effective, their application is likely to expand to other critical geographical regions with common characteristics such as dense commercial traffic and concentrated strategic infrastructure.

In this evolving environment, situational awareness alone is no longer sufficient to address the complexity of hybrid maritime threats. Traditional maritime monitoring systems like the Deimos MARVISION system are used to detect anomalies and breaches of standard behaviour patterns, but they rarely infer intent or focus on attribution, two key components for effective counterintelligence and deterrence. This gap has been identified as one of the most critical capabilities missing from western maritime security institutions today[11]. Without the ability to contextualise behaviour and identify potential actors or conflict points, anomaly detection alone leads to a purely reactive policy, which is easily defeated by oversaturation of the operational environment and deceptive action. By addressing these gaps in capability and challenging the plausible deniability narrative, grey-zone operations are made significantly harder to conduct as adversaries rely on delayed responses or lack of conclusive evidence to avoid escalation.

---

[8] Philippines accuses China's forces of harassing fisheries vessels in the South China Sea | AP News
[9] NATO Review - Fortifying the Baltic Sea - NATO's defence and deterrence strategy for hybrid threats
[10] All Together Now: China's Militia in 2025 | Asia Maritime Transparency Initiative
[11] Maintaining the UK's Intelligence Edge in the Grey Zone | Royal United Services Institute

In this context, it is important to distinguish between types of attribution, including technical, operational/strategic and political attribution. Technical attribution refers to the identification of threats based on specific evidence, with operational attribution referring to the analysis of behaviour patterns to determine the origin of the threat. Political attribution, which could be considered above the level of responsibility for the proposed model, involves considering the motivations and geopolitical interests of the actors involved. Because political attribution remains a policy decision, the model will focus on the first two, attempting to simplify an exceptionally complicated situation to allow for a better analysis. This distinction is crucial to the development of the system, as clarifying the boundary ensures that overreliance on automated outputs is prevented, making sure that the tool supports human judgement instead of replacing it.

2.2. Current monitoring approaches.

Current maritime monitoring tools rely on a combination of technological surveillance systems and manual analysis. These tools include AIS-based anomaly detection systems, which analyse vessel movement data to identify deviations from standard navigation patterns, manual Open Source Intelligence (OSINT) investigations relying on ownership records, sanctions lists and vessel location history or geofencing around strategic assets to sound alarms once a vessel has trespassed or intruded into a protected area. Despite their usefulness, as mentioned above, this approach remains reactive in nature, with most systems identifying isolated anomalies rather than analysing behavioural patterns over time, which is a task almost exclusively left to the analyst. This reactive posturing has been highlighted as one of the main structural limitations of current NATO and EU maritime awareness framework[12]. As a result, a large amount of adversarial activity is left unexamined, with even more threats remaining undetected as hybrid activity ramps up in the coming years. In particular, they rarely integrate predictive mechanisms and strategies like dynamic risk scoring or behavioural modelling, which means that current approaches

[12] *Beneath NATO's Radars: Unaddressed Threats to Subsea Cables | Strategic Technologies Blog | CSIS*. (n.d.). Retrieved 14 May 2026, from https://www.csis.org/blogs/strategic-technologies-blog/beneath-natos-radars-unaddressed-threats-subsea-cables

struggle to capture adaptive or covert behaviour, necessary in an environment that is ever changing. This leads to a lack of support for predictive or attribution-focused maritime security operations, resulting in a reactive policy rather than a proactive one.

The use of comparative case studies across different areas of operation is also crucial in today's interconnected space, where actors are no longer constrained to conducting hybrid operations inside their own area of influence. This can be seen through recent operations conducted in the Caribbean by nations like Russia, with ship routes ranging from the Persian Gulf to the Caribbean and the North Atlantic Ocean over a couple of days[13]. This highly mobile environment characteristic of modern-day grey zone threats also highlights the need to create comparative case studies between these different areas, allowing the algorithms to detect patterns that would likely go undetected otherwise. Another issue applications like this set to solve is the surplus of information of the modern intelligence collection environment, whose analysis is constrained by a shortage of collection platforms. By bringing information from multiple sources into a single analytical environment and highlighting the most urgent threats as classified by the predictive scoring, applications like these could work towards alleviating the "information overload" problem encountered by analysts[14], presenting a clear view of incidents that might otherwise go undetected in the stream of data collected from busy shipping lanes.

The current geopolitical climate is also beneficial to the development of similar projects, with commitments to increase defence expenditure for NATO members to 5%, with 1.5% of the committed GDP available to drive similar defence initiatives in dual use security-related technologies[15]. This allocation of resources highlights the need to develop strategic technology and data infrastructure to strengthen collective security, with capability building in this area being more crucial than ever due to the recent prevalence of hybrid threats.

Translating this political commitment into operational capability, however, depends on addressing several structural dependencies that extend beyond procurement and into the development of AI itself. The long-term sustainability of AI enabled intelligence capabilities within Europe are currently under scrutiny, with policy makers being concerned that the effort to develop effective models isn't succeeding. Current European

[13] Exclusive: US seizes Venezuela-linked, Russian-flagged oil tanker after weeks-long pursuit | Reuters
[14] Too Much Information: Ineffective Intelligence Collection
[15] NATO - Topic: Defence expenditures and NATO's 5% commitment

players in AI building like Mistral AI or Aleph Alpha have fallen behind the race to develop frontier AI models, not for a lack of qualified talent and researchers, but due to a lack of access to computational infrastructure. Frontier AI development now depends on access to large-scale data centres, advanced semiconductor supply chains, along with access to large scale energy generation capacity.

Importantly, Europe remains a critical actor within the AI development ecosystem, especially with companies such as ASML having a near monopolistic position in the production of extreme ultraviolet (EUV) lithography systems, critical in the manufacturing of advanced semiconductors used for AI platforms. However, a strong position in the semiconductor value chain does not automatically translate into leadership in frontier AI development. Transitioning from efficiency focused models to building large scale compute should also be a priority when development is becoming increasingly compute-constrained. Early in the development race, two competing assumptions arose along the lines of AI development, US based companies, with access to the latest semiconductor technology and large-scale private investment opted for a scaling oriented strategy. This approach was based on the assumption that increases in computational resources, training data and model parameters would yield unmatched improvements in capability. In contrast, developers operating under computational constraints, particularly after US semiconductor export controls on China, prioritised algorithmic efficiency, hoping to reduce dependency on computational infrastructure, allowing competitors with limited resources such as DeepSeek to remain competitive. While efficient training or machine learning techniques such as model distillation have reduced the computational cost of AI operations on an individual scale, these gains have been outpaced by the development of compute intensive methodologies like increased deployment, larger context windows or the reliance on autonomous agents[16]. This efficiency focused approach can also be considered to be flawed from its fundamental principles, showing characteristics seen at other times of increased technological acceleration such as the industrial revolution. This is clearly the case seen when referring to Jevons' paradox, originally developed as William Stanley Jevons, a victorian economist, observed that the development of more efficient steam engines did not reduce coal consumption but instead increased it as steam power became cheaper, leading to its widespread adoption. When

[16] https://doi.org/10.48550/arXiv.2511.15259

applied to AI development[17], one could argue that the development of efficient models would only drive the requirement for available compute, benefiting resource rich players, as has been the case recently.

Early investment is also a critical first step to close the capability gap, with increasing returns as compute capacity generated now not only improves current capability but accelerates future capability generation, especially before model self-training is fully implemented. Currently, limitations on self-training include the degradation of model quality and even model collapse, but research along this development vector has become a priority as realisation that high quality human generated data may become scarce in the future, making future available compute crucial in achieving exponential growth. An additional concern is that open-weight model development may not be appropriate in the current climate. While open weight models are fitting for EU goals with regards to data privacy and customisation (being self-hosted allows users to run models on private networks, ensuring total data privacy), it also includes some risks. In the first place, it is estimated that open-source models usually trail behind closed models by around 8 to 10 months[18], especially concerning in a field developing as fast as AI is. Furthermore, it also creates security concerns, especially tampering with model safety guardrails, bypassing safety features[19] or creating vulnerabilities when facing potentially malicious, unregulated foreign open weight models[20] due to the difficulty of extraterritorial enforcement. Finally, focusing on open-source model development also introduces risks regarding model use and download, eliminating the possibility of maintaining trade secrets and allowing for tech appropriation, essentially financing AI development for both legitimate users and competitors.

Decisions surrounding AI development often focus on algorithms, data and semiconductor manufacturing, while computational infrastructure often remains constrained by access to reliable and affordable energy. Large-scale AI development is becoming increasingly reliant on intensive energy use, creating a direct relationship

---

[17] Anh, D. T. (n.d.). *The Jevons Paradox in AI Infrastructure: Efficiency, Rebound, and the Limits of Cost Reduction*.

[18] How far behind are open models? — LessWrong
[19] Estimating Worst-Case Frontier Risks of Open-Weight LLMs
[20] Why we have six months to regulate the AI that Europe isn't building – CEPS

between energy generation capability and technological competitiveness. From a security perspective, energy infrastructure should not only be regarded as an economic asset but a strategic one, enabling future defence systems and intelligence capabilities. This also validates Lowery's power projection framework for future conflict, where natural proof of power in form of watts is a measure of capability and a form of conflict resolution, signalling computing power in this case. Furthermore, the same hybrid actors that target maritime infrastructure, undersea cables and critical infrastructure may increasingly view energy generation systems as attractive targets due to their importance in supporting digital economies and modern defence capabilities, with attacks on power stations[21] [22] and datacentres[23] already being framed as serious security threats. This convergence between energy policy and AI strategy requires policy makers to view energy security technological autonomy, both driving national security, as a single broad strategic framework rather than separate policy domains. As a result, investments in grid modernisation, secure energy generation and the protection of critical infrastructure must be prioritised to contribute to economic competitiveness, technological sovereignty and achieving national security objectives.

Ultimately, these vulnerabilities and structural flaws in the global compute race suggest that policy makers can no longer rely on a passive, efficiency first philosophy or depend on open weight model development. Instead, these challenges highlight the need to reassess current development priorities, placing a greater emphasis on the creation of secure energy generation capacity and the creation of sovereign large scale computational infrastructure.

2.3. Technological foundations behind maritime intelligence.

Over the past decades, advances in satellite communications, sensor technology and large-scale data processing have enabled the evolution of intelligence collection and data analysis. Modern maritime awareness systems combine cooperative tracking systems (that rely on active, compliant targets to share data, making them easy to interpret but vulnerable to interference) such as AIS, or non-cooperative (which don´t

[21] Hybrid Warfare and Cyber Effects in Energy Infrastructure | connections-qj.org
[22] Swedish power plant targeted by pro-Russian group in 2025, government says | Reuters
[23] Data Center Warfare: Defending the Key Terrain of AI Infrastructure - Modern War Institute

rely on target input) methods such as satellite imagery, radar and radio-frequency detection. These systems combined generate large volumes of data describing vessel movements in time and space, enabling the detection of irregular behaviour and threat identification. However, the expansion of intelligence capabilities has been further accelerated in the past decade through the development of machine learning and data analytics. By analysing vessel trajectories, behavioural characteristics and ownership networks among others, these systems can now support predictive risk assessments and intelligence gathering in complex environments. The following sections review the main technological aspects that uphold current monitoring systems and their role in maritime intelligence collection.

2.3.1. Automatic Identification System (AIS).

The Automatic Identification System (AIS) is the main cooperative tracking system used in modern maritime surveillance, relying on GNSS (Global Navigation Satellite System) to transmit any ship´s position as well as their speed and course. Developed by the International Maritime Organisation (IMO), the International Telecommunication Union (ITU) and the International Electrotechnical Commission (IEC) under the SOLAS agreement (Safety of Life at Sea) and made compulsory in December 2002[24] for all passenger vessels and international voyaging ships with a gross tonnage higher than 300.

AIS works by transmitting a stream of information to nearby stations through two VHF radio channels, at 161.975 and 162.025 MHz as defined by the ITU. From a technical standpoint, AIS communications take place using a Self-Organized Time Division Multiple Access (SOTDMA) protocol, which divides time into discrete slots (2250 per minute, per channel). Each data transmission includes the TDMA slot reference being used by the transmitting station, building a “map” of used slots in the region. This allows multiple AIS systems to share the same frequency by selecting transmission slots based on local traffic density, modifying its transmission slot if required without needing a central station to organise the transmissions. Another advantage of this system is that, in case of oversaturation

[24] Automatic Identification System (AIS) - Great Lakes St. Lawrence Seaway System

in a crowded environment where all TDMA slots are occupied, the AIS transmissions will occupy the slots with lower received signal strength being used by the most distant mobile stations, reducing the range of the AIS system but ensuring communication is maintained. Transmission rate is also variable as a function of vessel dynamics, with fast moving ships transmitting data packets every 2-10 seconds with ships at anchor updating their situation every couple of minutes, balancing bandwidth efficiency with situational awareness requirements automatically.

This information can be split into three categories, static, navigational and voyage specific. Static information refers to the fixed attributes of the transmitting ship, including the vessels Maritime Mobile Service Identity number (MMSI), International Maritime Organization number (IMO), call sign, name, length and beam and type of ship. This information is crucial for ship identification, but does not provide operational details, making it hard to identify threats solely based on Navigational information, excluding ships already included in sanctions lists or previously identified as possible threats. Vessel characteristics that are voyage-specific are also included in the AIS transmissions, including the ship's draught, route plan (with destination and ETA), as well as if the vessel carries hazardous cargo. Lastly, navigational information includes vessel speed, course, heading, navigational status and rate of turn parameters, which could be used to infer intent through behavioural analysis.

While AIS has proven to be an irreplaceable tool for ensuring ship to ship communication as a navigation aide, as well as a source of data collection for purposes such as accident investigation, it is also vulnerable to attacks or misuse in the modern maritime environment. In first place, one of the limiting factors for AIS is the lack of effective methods for collecting, organising and interpreting the data. The lack of a standard, open access, AIS benchmark database makes developing algorithms working on AIS data more complicated. The lack of a standard database, combined with the vast amount of data streamed in real time is also an important contribution to the "information overload" problem which this project aims to alleviate. However, it is important to highlight some broader issues which are crucial to understand when using AIS as a data source, which will be

explored in the following section, while giving some context of their utility for hybrid operations.

2.3.1.1. AIS Hijacking/Spoofing:

"Spoofing" refers to the act of transmitting false or misleading AIS data for concealment. In this type of attack, actors, either working from the ship itself or through cyberattacks/radio frequency modification, assign erroneous static, navigational, or voyage-specific information to the desired ship to create the illusion of standard behaviour. In other attacks, spoofing might be used to harass ships, creating navigation issues in congested waterways or even impeding search and rescue missions. This has been the prevalent mechanism for AIS interference in the case of current targeted hybrid operations. However, due to its prevalence in the early development of hybrid maritime operations, enforcement agencies prioritised the monitoring of ship spoofing, leading to a high incident-sanction rate. As found by the analytics platform kpler, 80.1 percent of ships caught spoofing in their sample of 261 ships were sanctioned within a year, with 61 vessels being sanctioned within 3-6 months and 59 vessels being sanctioned within 6-9 months[25]. Despite the high enforcement rate and becoming a telltale sign for identifying threats, ship spoofing will continue to be one of the main tools for adversaries to conceal their real activity, degrading situational awareness and enabling plausible deniability, masking covert operations.

In many cases, spoofing events are recorded on AIS as abnormal positional behaviour, with several of the behavioural anomalies being coordinate jumping or abnormal changes in vessel speed. In most cases when the spoofing objective is the denial of service, the data obtained shows three clear signs, relocation to multiple targets (often in a circular pattern), relocation to multiple static targets presented as point locations, or relocation of a smaller operational area to a single static target. These

[25] AIS spoofing: The fast track to sanctions | Kpler - Nov 10, 2025

recurring behavioural patterns demonstrate how, while designed to obscure true vessel movement, these anomalies can be systematically detected. The nature of these spoofing attacks also means that analysis of the signals is oriented to interpretation using machine learning algorithms and behavioural analysis, as the objective shifts from verifying individual data points to understanding vessel movements over time.

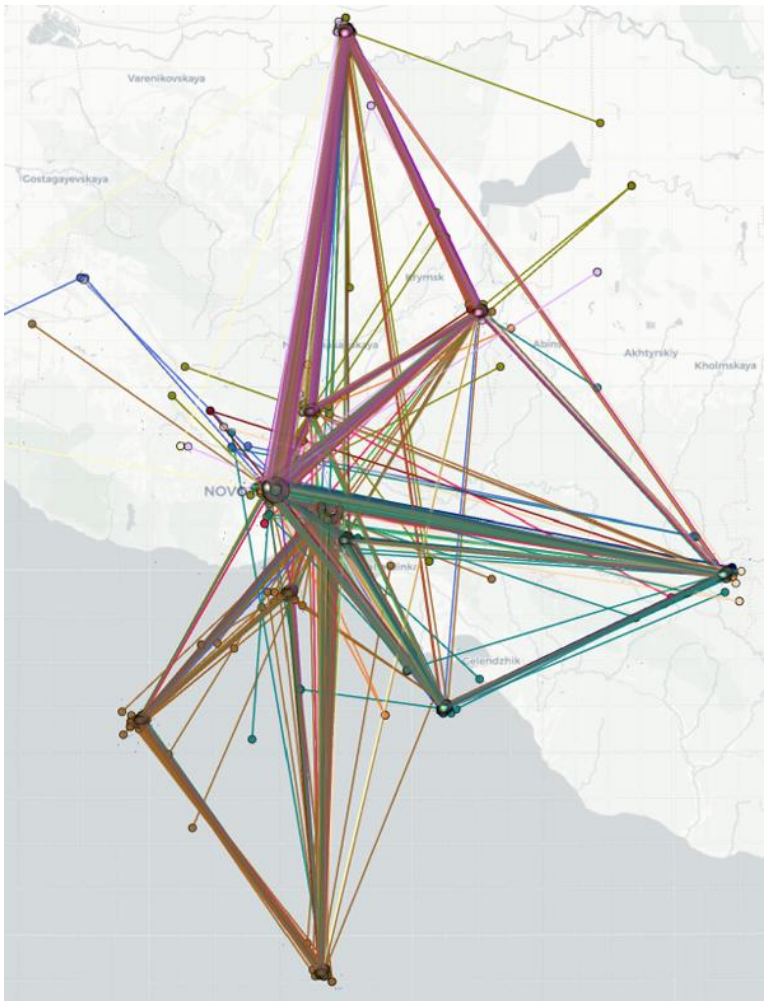

Fig. 1. FPSEs (Final Potential Spoofing Events) around Novorossiysk at the Black Sea, recorded on 21-23 January 2026. Source: SeaSpoofFinder -- Potential GNSS Spoofing Event Detection Using AIS

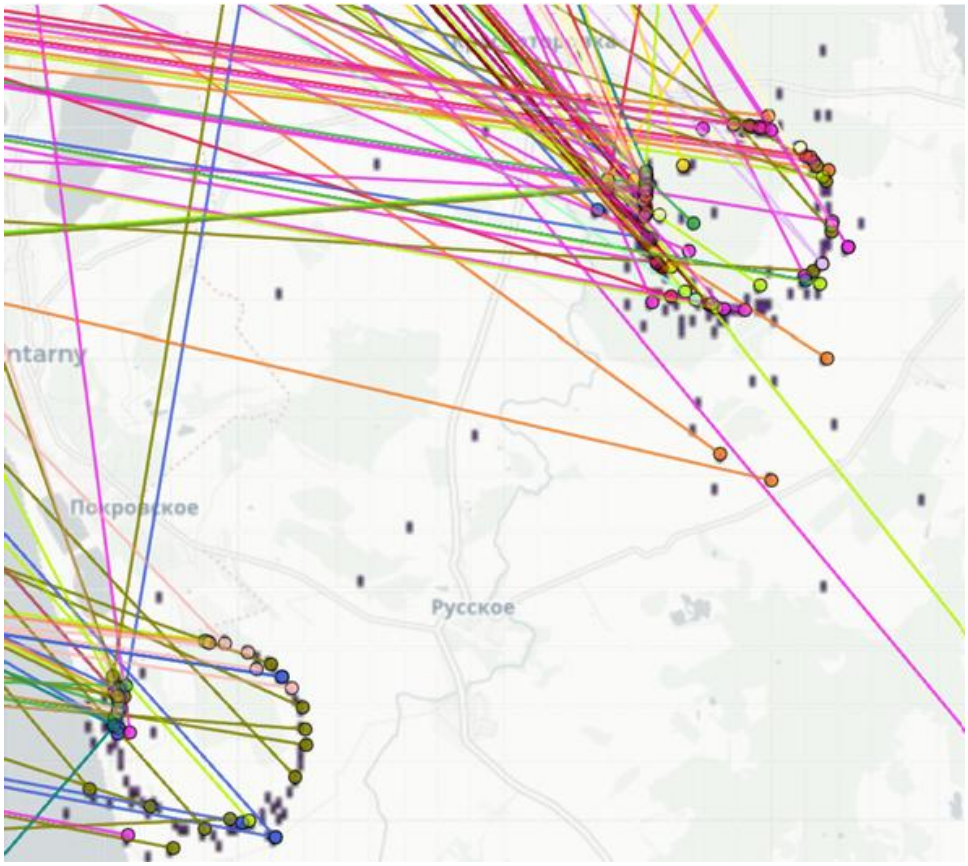

Fig. 2. FPSEs in circles in Kaliningrad, recorded on 18-20 January 2026. Source: SeaSpoofFinder -- Potential GNSS Spoofing Event Detection Using AIS

2.3.1.2. Aids to Navigation (AtoN) Spoofing:

In a similar manner to ship information, vessel navigation aids, such as beacons or hazard warnings, can be manipulated to show false information, leading ships into dangerous situations or into conducting unwanted manoeuvres. In the context of grey zone operations, this type of manipulation can be easily used to disrupt traffic across choke points or disrupt port access, creating economic bottlenecks. For example, the "closing" of ports or energy terminals could trigger insurance spikes and force rerouting, redirecting shipping into unsafe waters. Even if this manipulation mechanism could be used to target individual ships, causing accidents or unsafe situations, the main use case is the disruption of maritime infrastructure for an area of operations through economic coercion rather than kinetic action.

2.3.1.3. Collision Spoofing.

As opposed to AtoN Spoofing, Collision Spoofing is a targeted spoofing mechanism which has the primary purpose of creating kinetic disruption in which attackers deliberately manipulate the Closest Point of Approach (CPA) feature on the vessel's AIS system to trigger collision alerts. CPA, one of the central features of AIS systems, is designed to calculate the minimum distance between two ships, taking into account velocity vectors to analyse collision probabilities, especially useful in congested maritime traffic. By manipulating the CPA system, attackers aim to cause abrupt heading changes, creating "accidental" incidents in which vessels are forced into dangerous manoeuvres, potentially colliding with other vessels or running aground. Like the other spoofing mechanisms, this is especially relevant to strategic chokepoints with dense commercial traffic, in which the probability of causing meaningful damage through collisions rises exponentially.

2.3.1.4. AIS-SART Spoofing.

Another of the main uses for AIS as declared by the SOLAS agreement is to facilitate search and rescue operations, detecting vessels in distress and updating their voyage specific status, alerting the authorities and providing reliable location data. In this scenario, an attack would involve the manipulation of the vessels' Search And Rescue Transponder (SART) to simulate a fake emergency.

As stated in Article 98 of the United Nations Convention on the Law of the Sea[26] "Every State shall require the master of a ship flying its flag, in so far as he can do so without serious danger to the ship, the crew or the passengers:

a) to render assistance to any person found at sea in danger of being lost;

[26] INTERNATIONAL TRIBUNAL FOR THE LAW OF THE SEA

b) to proceed with all possible speed to the rescue of persons in distress, if informed of their need of assistance, in so far as such action may reasonably be expected of him;
c) after a collision, to render assistance to the other ship, its crew and its passengers and, where possible, to inform the other ship of the name of his own ship, its port of registry and the nearest port at which it will call"

By creating fake emergencies, adversaries can exploit AIS-SART spoofing in multiple ways, testing and gathering intelligence on response times and procedures, luring vessels into ambush zones or manipulating naval and coast guard deployment, pulling assets from sensitive areas and creating "power vacuums" that can be easily exploited.

2.3.1.5. Weather Forecasting Manipulation:

AIS integrates a weather reporting system in which dynamic environmental data pertaining to weather, currents or other relevant climate conditions, is communicated between meteorological stations and vessels at sea. Again, an attack on the system could be used to manipulate traffic by emitting false alerts, rerouting traffic and disrupting supply chains with difficult attribution.

2.3.1.6. Availability Disruption Threats:

Availability disruption refers to the type of attacks designed to disable AIS systems for targeted ships or in a localised area. By creating localised denied areas or "blackout" zones, actors aim to disrupt shipping, compromising its safety or enable covert operations. There are three main cases of availability disruption threats seen in today's maritime environment, with the first one being Slot Starvation. This type of attack involves impersonating the relevant maritime authority to block AIS transmission space in order to create a large-scale regional blackout. An example of a selective denial attack rather than widespread area denial would be a Frequency Hopping attack, in which targeted vessels are forced

into unusable frequencies by an impersonated maritime authority. By protocol, the targeted AIS station is required to maintain the new frequency even if the system is rebooted, making the attack difficult to solve. This attack can also have a geographic component, in which the aggressor programs the targeted system into switching frequencies once the ship reaches a designated area. Lastly, the Timing Attack consists on preventing transponders from communicating their position by instructing the AIS system to repeatedly delay it´s transmissions, making a vessel disappear from AIS based tracking systems. In another version of this attack, attackers can overload AIS information systems through DDOS-like approaches by requesting transmissions at a high enough rate.

As a cooperative tracking mechanism, AIS is especially vulnerable to defiant behaviour through intentional deactivation. Despite being a compulsory system for safety purposes under the International Maritime Organisation, AIS can be manually disabled by vessel operators. This issue lies in the fundamentals of AIS, as it was conceived originally as a safety system instead of a surveillance or enforcement mechanism, relying on voluntary compliance as a common structural limitation of cooperative surveillance architectures.

In operational contexts, disabling AIS is a telltale sign of illicit activities, including sanctions evasion, illegal fishing or ship to ship transfer of sanctioned or illicit goods, allowing for the interpretation of the behavioural anomaly to derive the missing data. This can be done by analysing the manoeuvres in a multi-sensor context, integrating non-cooperative detection methods such as Synthetic Aperture Radar satellite imagery. This transforms non-cooperative behaviour into an intelligence signal, inferring potential intent based on spatial and temporal patterns surrounding the incident. As a consequence, AIS should not be understood as a means to achieve complete awareness of the maritime domain, but as a layer within an integrated intelligence architecture, with limitations that require external sensing capabilities and analytical methods to ensure resilience against adversarial manipulation.

2.3.2. Satellite Surveillance:

Satellite surveillance, as opposed to AIS, is the main non-cooperative technology backing maritime intelligence in the current era, enabling large scale and persistent monitoring of shipping activity. Unlike AIS, which relies on self-reported data, satellite systems obtain information physically from electromagnetic radiation, either passively through optical or infrared sensing, or actively through methods like radar.

Relevant satellite-based intelligence collection systems can be sepated mainly into three types, Electro-Optical imaging platforms, Synthetic Aperture Radar and Radio-Frequency platforms.

Electro-Optical satellite systems use passive optical systems to capture imagery along the visible and infrared spectrum. They are characterised by providing high spatial resolution, although limited by atmospheric conditions such as cloud cover, haze or darkness. This makes them suited to tasked imaging assignments rather than continuous monitoring, even if they prove to be great systems for image analysis and visual confirmation. EO systems include Keyhole (KH) systems, including the first CORONA satellites built by Lockheed Martin under contract to the CIA. Successive KH platforms developed quickly during the cold war, including the GAMBIT or HEXAGON series, and are now managed by the NRO.

On the other hand, Synthetic Aperture Radar (SAR) satellites use active sensing by illuminating the target with microwave radar and measuring the reflected signal. The main advantage that suits them to maritime intelligence is their ability to detect vessels regardless of atmospheric conditions, being able to obtain data through cloud cover or in the dark, allowing for continuous monitoring. As the name indicates, SAR achieves high resolution by simulating an antenna through its forward movement, doing so by illuminating the target from different positions along its trajectory, effectively synthesising a long antenna. Each pulse returns a slightly different view of the target scene, which the system then reconstructs by

using doppler shifts and phase differences, overcoming the traditional limitation for resolution of antenna length (longer antennae produce narrower beams and finer resolution). SAR satellite concepts first went under development with the NRO-driven QUILL program, with a launch in 1964, with more recent programs being outsourced to providers such as ICEYE US, contracted to provide persistent monitoring. It is the ability to provide persistent monitoring in all atmospheric conditions that make these satellites uniquely fitted to operating in constellations, as they can provide imagery at any time, making the large investment justifiable. By doing so, these systems are able to reduce revisit time drastically, enabling persistent monitoring and improving the imagery result through obtaining data at different look angles, reducing shadowing and improving the ability to detect vessel wake among other characteristics.

While SAR and EO provide imagery that can be analysed, RF sensing intercepts electronic signatures emitted by the observed target, contributing to forming a complete picture through Signals Intelligence (SIGINT) or Electronic Intelligence (ELINT). They operate passively without illuminating the target and can detect vessels by matching their unique RF signatures to collected databases, even if AIS is spoofed or disabled. These satellites can also operate in constellations with different satellites bearing different sensing instruments, helping to obtain a complete picture through data fusion. Operating in constellations also enables these satellites to geolocate targets through algorithms like Time Difference of Arrival (TDOA), which measures the time delay of a signal arriving at two separate receivers or through Frequency Difference of Arrival (FDOA), which measures the difference in doppler shift between two receivers to estimate a Line of position along which the target lays.

These systems allow for intelligence collection over denied areas without direct presence, and, in many cases, without the target being alerted to their surveillance, making it more difficult to hide adversarial behaviour. However, satellite surveillance is defined by three main constraints, orbital mechanics, sensor capabilities and data processing architectures, which will be explored in the following sections.

2.3.2.1. Orbital Mechanics and Coverage Constraints:

Satellite performance and the role it fulfils is heavily dependent on the orbit it follows, as it governs the coverage geometry, temporal resolution and sensing conditions, which determine the sensing payload that can be carried to optimise intelligence gathering.

Most maritime surveillance satellites operate in Low Earth Orbit (LEO), at an altitude between 180 and 2000 km. At these altitudes, the orbital period is usually around 90-110 minutes, allowing multiple passes per day over different regions of the Earth and a higher revisit rate (or how frequently it can observe the same location) than other orbits. An example of maritime surveillance satellites in LEO includes the Naval Ocean Surveillance System (NOSS), codenamed PARCAE (1st generation), RANGER (2nd gen.) or INTRUDER (3rd gen.), a U.S. Navy ELINT platform first launched in 1976 as a system designed to geolocate Soviet Navy assets during the Cold War. This system began to be operated by the National Reconnaissance Office (NRO) as the agency took over Satellite Surveillance programmes overtly after its declassification in 1992, with its most recent launch taking place in March 2025 .This system had a roughly circular orbit, with an apogee (furthest point from Earth in orbit) and perigee (closest point from Earth in orbit) of roughly 1100 km and an inclination of 63.4º, known as the critical inclination. It is important to note that this orbit is characteristic of surveillance satellites as the inclination allows the satellite to remain in a stable orbit, cancelling out the shift of the perigee the equator bulge causes on regular orbits.

Another reoccurring orbit for surveillance platforms are sun-synchronous orbits, polar orbits with an inclination of roughly 98º and a LEO of 600-800 km that match the earth's rotation around the sun. By keeping local time at the moment of passing constant, trends can be monitored more consistently, since light and shadow conditions are comparable between observations. This allows analysts to easily compare pictures and identify

differences indicating developments, also making it a crucial parameter for change detection algorithms.

The selection of LEO is driven by two competing requirements, spatial resolution or Ground Sampling Distance (GSD), which determines the detail of the captured imagery, and swath width (W), which determines the observable ground strip.

$$GSD = \frac{H \times P}{f}$$

With H being the orbital height, P the detector pixel size and f being the focal length of the sensor.

$$W = 2H \times \tan(\alpha)$$

With α being the sensor field of view.
As can be derived from these formulas, selecting the orbit is a trade off between wide area surveillance or fine detail imaging, with satellites being designed for a variety of coverage roles. However, higher orbits such as Medium Earth Orbit or Geostationary orbits are generally unsuitable for higher resolution maritime surveillance due to the degradation in spatial resolution and signal strength over the longer distances.

Another critical parameter for maritime surveillance satellites is the revisit time, briefly mentioned above. This parameter determines the interval between successive observations for a given target and is governed by the orbital period, Earth's rotation and the sensor swath width. A single satellite in LEO has a typical revisit time of 1-3 days, depending on its latitude and sensor characteristics. This temporal coverage gap is a vulnerability that can be exploited for covert activities such as ship-to-ship transfers of entry and exit from restricted zones. This limitation historically constrained space reconnaissance to "snapshot intelligence" rather than continuous monitoring. However, to mitigate this issue, modern systems

use satellite constellations to distribute satellites among different orbital planes, maximising geographic coverage, appropriate temporal spacing of observations (through satellite phasing) and reducing revisit time by increasing the number of satellites, allowing for persistent surveillance instead of predictable passes (frequently called deterministic observation).

Another way of increasing effective coverage for satellite imagery is through adjusting the sensor's look angle, only available through developments in signal processing. This allows an increased revisit frequency as well as the ability to target specific geographical areas dynamically. However, this comes at a cost, reducing the spatial resolution and increasing the geometric distortion in imagery.

2.3.2.2. Implications for Maritime Intelligence Systems:

By understanding satellite coverage and designing orbital parameters for effective monitoring, satellite surveillance systems become crucial for maritime surveillance. However, orbital mechanics impose a set of constraints that systems must try and expand in order to become effective against the expansion of hybrid operations. As seen above, these constraints can be mainly summarised in three points.

Firstly, satellite coverage is discrete rather than continuous, which is why maritime intelligence satellite architectures must integrate multi-orbit constellations instead of relying on a single system with standard, predictable characteristics. This introduces the temporal uncertainty problem, in which vessel activity between observation windows cannot be observed but must be inferred instead. As a result, the reconstruction of activities and trajectories followed by the vessel becomes a probabilistic process, with a larger uncertainty in the estimated vessel state as the revisit time increases. This means that the possible area in which the vessel is located expands over time, forming an uncertainty region that grows between observations. One possibility to represent this issue would be to use a covariance matrix for the vessel state, populated by the uncertainties

in X position, Y position and velocity. In this case, the total uncertainty can be represented by the trace of the uncertainty covariance matrix, giving a measure of the total uncertainty in the estimated vessel position. This can be represented for a given scenario by modifying the revisit parameters as the independent variable in the case study, as demonstrated below.

In order to get real data for the derivation, the sanctioned, Iranian-flagged bulk carrier "ARVIN" (IMO 9193202) has been selected. At the time of study, the vessel was shown to be underway using engine at the Riau Archipelago between the port of Bandar Abbas, Iran and the port of Dafeng, China, with a speed over ground (SOG) of 9.3 knots, a course over ground (COG) of 47º, as well as a rate of turn (ROT) of 0º/min.

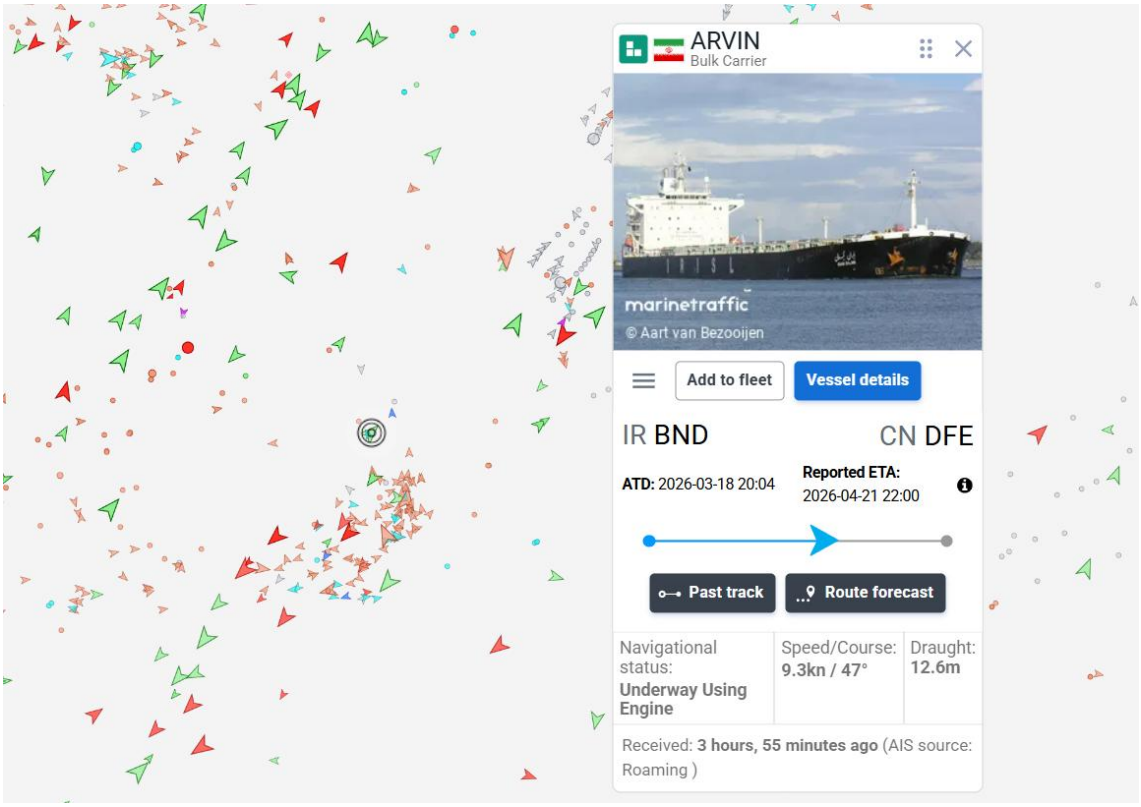

Fig. 3 Image showing the AIS status for the vessel "ARVIN" on 11/04/2026.

After converting the data into SI units, it can be inputted into the state vector containing the data from the AIS observation, including position (x and y), and velocity (Vx and Vy).

$$State\ vector = [x \quad y \quad v_x \quad v_y]$$
$$V\ (m/s) = SOG\ \times 0.5144$$
$$V_x\ (m/s) = V\cos(\theta)$$
$$V_y\ (m/s) = V\sin(\theta)$$

It is important to note that, while x and y could be inputted as location coordinates, the initial position of the vessel is taken as (0,0) as a reference point for this exercise.

Matlab code:

```
SOG_kts  = 9.3;                    % Speed over ground [knots]
COG_deg  = 47;                     % Course over ground [deg]
v        = SOG_kts * 0.5144;       % Conversion to SI units
th       = deg2rad(COG_deg);

vx = v * cos(th);
vy = v * sin(th);
```

The values for σ_x and σ_y, associated with the position uncertainty from the AIS readout data, which originates from GPS error or receiver quality, will be taken as 5m. This is justified by taking the values for GNSS accuracy for maritime navigation[27], which take horizontal accuracies for C/A-code GPS to be around 5-10 m (1σ). When differential corrections (DGPS/SBAS) are available, the accuracy improves to 1-3 m (1σ), however, "ARVIN" AIS transponder class is class A, justifying the selection of 5 m.

The velocity uncertainty components can be obtained by applying the first-order uncertainty propagation rule as stated from Bar-Shalom et al, *Estimation with Applications to Tracking and Navigation*, basis for most

[27] GPS Performances - Navipedia

of the following formulas and the Kalman filter estimation process methodology for this application.

$$\sigma_y^2 = \sum_i \left(\frac{\partial f}{\partial x_i}\sigma_{x_i}\right)^2$$

The partial derivatives can be obtained from the state vector as follows:

$$\frac{\partial v_x}{\partial v} = cos\theta$$

$$\frac{\partial v_x}{\partial \theta} = -v\ sin\theta$$

$$\frac{\partial v_y}{\partial v} = sin\theta$$

$$\frac{\partial v_y}{\partial \theta} = v\ cos\theta$$

Resulting in the following velocity component uncertainties:

$$\sigma_{v_x}^2 = (cos\theta\ \sigma_v)^2 + (v\ sin\theta\ \sigma_\theta)^2$$

$$\sigma_{v_y}^2 = (sin\theta\ \sigma_v)^2 + (v\ cos\theta\ \sigma_\theta)^2$$

Using the uncertainty parameters for both the position and velocity uncertainties, the initial covariance matrix (Σ_0) can be constructed, representing the uncertainty about the vessel's state at the start of the estimation process.

$$\Sigma_0 = \begin{bmatrix} \sigma_x^2 & 0 & 0 & 0 \\ 0 & \sigma_y^2 & 0 & 0 \\ 0 & 0 & \sigma_{v_x}^2 & 0 \\ 0 & 0 & 0 & \sigma_{v_y}^2 \end{bmatrix}$$

It is important to note that the covariance values that are not on the diagonal term will be 0 as the AIS errors in position and velocity are independent.

Matlab code:

```
sig_pos   = 5;                      % AIS horizontal accuracy
sig_v     = 0.051;                  % Speed uncertainty [m/s]
sig_th    = deg2rad(2);             % Course uncertainty [rad]

sig_vx = sqrt((cos(th)*sig_v)^2 + (v*sin(th)*sig_th)^2);
sig_vy = sqrt((sin(th)*sig_v)^2 + (v*cos(th)*sig_th)^2);

P0 = diag([sig_pos^2, sig_pos^2, sig_vx^2, sig_vy^2]);
```

The vessel's motion will be modelled using a constant velocity model, assuming that the position evolves linearly with velocity and that the velocity remains constant between satellite observations.

Therefore, the equations for vessel position and velocity will be as follows:

$$x = x_0 + v_x t$$

$$y = y_0 + v_y t$$

$$v_X = v_{x0}$$

$$v_y = v_{y0}$$

In matrix form:

$$\begin{bmatrix} x \\ y \\ v_x \\ v_y \end{bmatrix} = \begin{bmatrix} 1 & 0 & t & 0 \\ 0 & 1 & 0 & t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_0 \\ y_0 \\ v_{x0} \\ v_{y0} \end{bmatrix}$$

With $F(t) = \begin{bmatrix} 1 & 0 & t & 0 \\ 0 & 1 & 0 & t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$, being the state transition matrix, that forms the deterministic part of the Kalman filter prediction mechanism, defining how the state evolves between satellite revisits.

Matlab code:

```
F = @(d) [1 0 d 0; 0 1 0 d; 0 0 1 0; 0 0 0 1];
```

The probabilistic aspect of the model is introduced by the process noise matrix (Q), that represents the uncertainty about the vessel's motion

between satellite revisits, quantifying how much the vessel can deviate from the constant velocity model. This uncertainty comes from unmodeled dynamics such as small accelerations, rudder corrections, engine fluctuations or course alterations caused by waves or currents.

The matrix, taken from *Estimation with Applications to Tracking and Navigation*, is composed by the following terms:

- $\frac{\Delta t^3}{3}$, the position variance derived from the integration of the acceleration (which in turn is the integration of the velocity).
- $\frac{\Delta t^2}{2}$, the position – velocity covariance, linking velocity uncertainty to position uncertainty.
- Δt, the velocity variance due to acceleration noise.
- q, the process noise intensity, representing the level of unpredictability.

It is important to note that the q value is not a set value and instead depends on the type of vessel being tracked. However, it can be approximated using the maximum acceleration as shown in the paper titled "Estimating Optimal Tracking Filter Performance for Manned Manoeuvring Targets" written by Robert A. Singer, obtaining an approximate value for large vessels.

$$q = \sigma_a^2$$

$$\sigma_a = \frac{a_{max}}{\sqrt{3}}$$

With the average maximum acceleration of a large tanker or cargo ship being around 0.025 m/s^2.

$$Q(t) = q\begin{bmatrix} \frac{t^3}{3} & 0 & \frac{t^2}{2} & 0 \\ 0 & \frac{t^3}{3} & 0 & \frac{t^2}{2} \\ \frac{t^2}{2} & 0 & t & 0 \\ 0 & \frac{t^2}{2} & 0 & t \end{bmatrix}$$

Matlab code:

```
a_max = 0.025;
q     = (a_max / sqrt(3))^2;
Q = @(d) q * [d^3/3  0      d^2/2  0    ;
              0      d^3/3  0      d^2/2;
              d^2/2  0      d      0    ;
              0      d^2/2  0      d    ];
```

Another important factor to take into consideration is the output of satellite measurements, which does not measure velocity, measuring the vessels position exclusively. This means that the measurement vector (z) will include the first two components of the state vector as well as the noise associated to the reading compared to the vessel's state.

$$z = \begin{bmatrix} x \\ y \end{bmatrix} + noise.$$

$$z_k(at\ time\ step\ k) = Hx_k + v_k$$

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

With H being the measurement matrix, indicating what components are included in the satellite's measurements. Because the satellite measurements also include errors, we must include a measurement noise matrix (R), quantifying the uncertainty of the satellite's position measurements, following Bar Shalom's tracking process.

$$R = \begin{bmatrix} \sigma_{z_x^2} & 0 \\ 0 & \sigma_{z_y^2} \end{bmatrix}$$

Matlab code:

```
R = diag([10^2, 10^2]);
H = [1 0 0 0; 0 1 0 0];
```

In this case, the variance of the noise affecting the satellite's position measurements will be taken to be 10 m, an average value for satellites such as RADARSAT-2[28] , and will be taken to be equal in both x and y directions.

Once all of the data has been arranged into their corresponding tables, the evolution of uncertainty between satellite observations can be modelled using the covariance propagation equation.

$$\Sigma(t) = F(t)\,\Sigma_0\,F(t)^T + Q(t)$$

The first term, $F(t)\,\Sigma_0\,F(t)^T$ describes how existing uncertainty develops through time, reflecting how errors in position and velocity grow as time progresses. The second term $Q(t)$, adds the process noise matrix, introducing the uncertainty created by unmodeled factors such as currents and waves. Together, these terms represent the total uncertainty at a time t as the sum of past uncertainty being carried forward and the new uncertainty added during the "dark" window.

Once the satellite revisits the area, obtaining data from the vessel once again, the Kalman filter is applied to correct the predicted state and reduce the uncertainty. However, this new observation does not reset the uncertainty to 0, as the data retrieved from the satellite includes error terms included in the R noise covariance matrix, as shown below.

$$K = \Sigma H^T (H\Sigma H^T + R)^{-1}$$

The value for the Kalman gain is then applied to the new satellite measurement, determining the weight assigned to the new measurement relative to the prediction (a large R reflecting noisy satellite measurements

[28] Geolocation Accuracy Validation of High-Resolution SAR Satellite Images Based on the Xianning Validation Field

will result in a small correction, while a small R, reflecting accurate measurements will result in large corrections). The updated covariance is then calculated as $\Sigma(I - KH)\Sigma$.

Matlab code:

```
%% SIMULATION PARAMETERS
dt      = 60;                        % Step [s]
T_total = 3600 * 12;                 % 12 hour simulation
t       = 0 : dt : T_total;
N       = length(t);

revisits = [3600, 3*3600, 5*3600]; % 1 h, 3 h, 5 h
labels   = {'1h Revisit', '3h Revisit', '5h Revisit'};
colours  = {'b', 'g--', 'r:'};


traces   = zeros(N, 3);              % RMS positional uncertainty
P_bubble = cell(3,1);                % Covariance just before each
revisit (max uncertainty)

for s = 1:3
    P = P0;
    rev = revisits(s);
    P_pre = P0;                      % Stores the last covariance
values before the update

    for k = 1:N
        tk = t(k);

        P = F(dt) * P * F(dt)' + Q(dt);

        %
        if mod(tk, rev) == 0 && tk > 0
            P_pre = P;               % Gets the covariance values
before the update pre-update uncertainty
            K = P * H' / (H * P * H' + R);
            P = (eye(4) - K*H) * P;
        end

        traces(k, s) = sqrt(trace(P(1:2, 1:2))); ;              %
RMS position uncertainty

    end
    P_bubble{s} = P_pre;
end
```

Two useful graphs can be obtained from the resulting values to visualise the effect of revisit time on positional uncertainty. The first case uses the trace of the covariance matrix to represent the evolution of the position

error covariance. Between the satellite visualisation windows, the filter relies solely on the motion model, showing two clear phases. After the vessel has been observed, uncertainty increases following a cubic polynomial if the process noise term dominates the covariance propagation equation, and according to a quadratic equation when the model-propagated term dominates. The second phase is shown during the simulated revisit period, where the Kalman update reduces uncertainty. The magnitude of this drop is determined by the measurement noise R. An increased frequency of the Kalman update contributes to keeping uncertainty errors low, while a higher revisit time shows the accumulation of uncertainty during long blackout periods as seen in the graph below.

Matlab code:

```
%% REVISIT FREQUENCY VS. POSITIONAL UNCERTAINTY
figure(1); hold on; grid on;
for s = 1:3
    plot(t/3600, traces(:,s), colours{s}, 'LineWidth', 1.8);
end
xlabel('Time [hours]');
ylabel('RMS Position Uncertainty [m]');
title('Revisit Frequency vs. Positional Uncertainty');
legend(labels, 'Location', 'northwest');
```

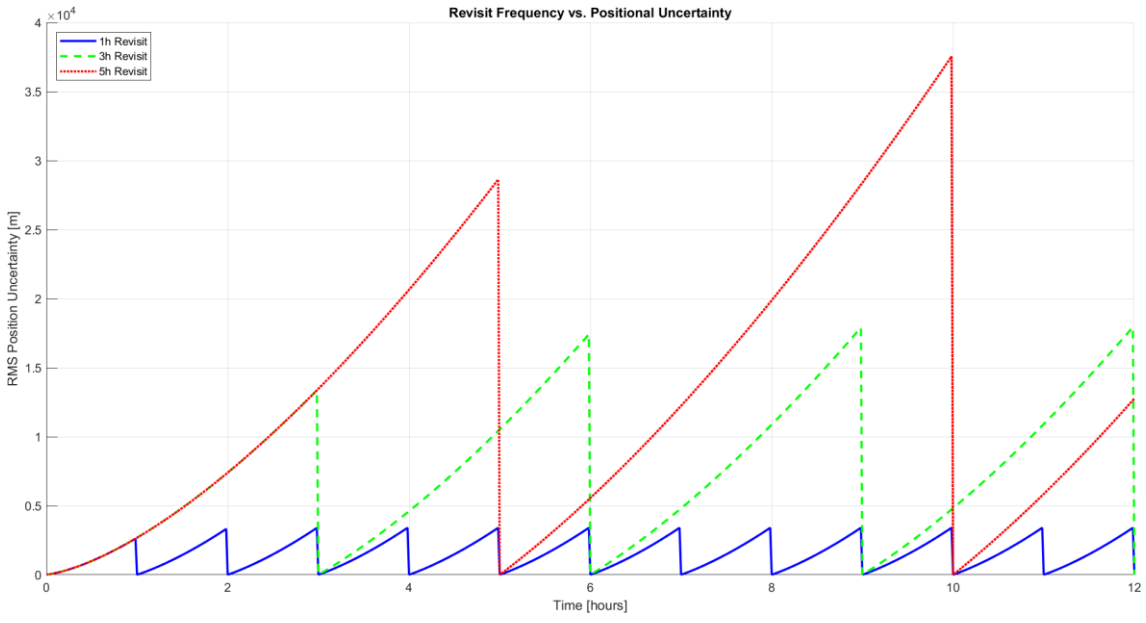

Fig. 4. Plot showing the increasing accumulated positional uncertainty as revisit frequency increases.

The second graph plots the uncertainty regions for the vessels position at the moment of maximum uncertainty, just before a satellite revisit. The regions are computed from the eigenvalues and eigenvectors of the position covariance matrix. An important note is that, for this case study, the ellipses appear circular due to the simulation constraints, stating that the process noise Q, the initial position uncertainty and the measurement noise R are symmetrical in x and y. However, the uncertainty region is typically an elongated ellipse aligned with the vessels direction of motion, as the velocity uncertainty projects into position uncertainty more strongly along the direction of travel and environmental disturbances are rarely isotropic, creating stronger disturbances and leading to a higher uncertainty along specific axes. Despite this modelling variations, the comparison between uncertainty regions for different revisit times is clearly visible in the graph below.

Matlab code:

```
%% MAXIMUM UNCERTAINTY REGIONS BY REVISIT INTERVAL

figure(2); hold on; grid on; axis equal;
th_c = linspace(0, 2*pi, 200);
unit = [cos(th_c); sin(th_c)];

for s = 1:3
    [V, D] = eig(P_bubble{s}(1:2, 1:2));
    ell    = V * (2 * sqrt(D)) * unit;   % 2σ ellipse
    plot(ell(1,:), ell(2,:), colours{s}, 'LineWidth', 2);
end
xlabel('Easting Error [m]');
ylabel('Northing Error [m]');
title('Maximum Uncertainty Regions by Revisit Interval');
legend(labels, 'Location', 'best');
```

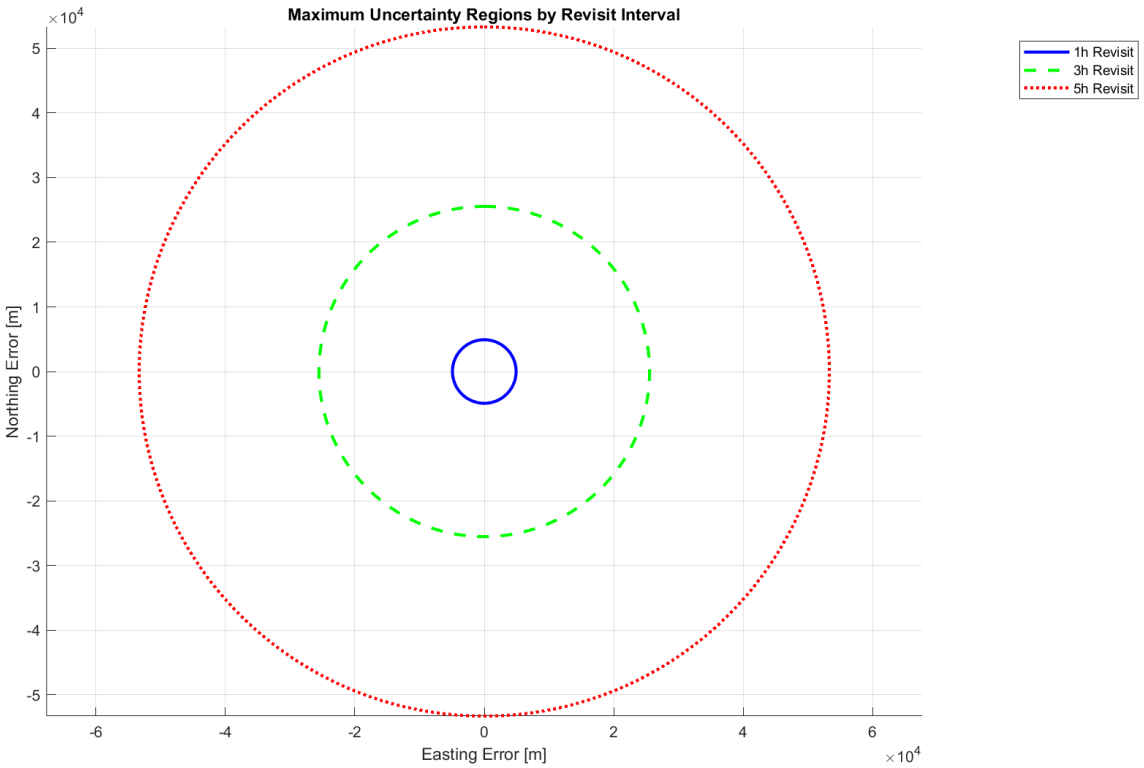

Fig. 5. Plot showing the increase in maximum uncertainty regions as a function of the revisit interval.

The two graphs shown above help to illustrate the impact of satellite revisit frequency on reliable vessel monitoring, with uncertainty accumulating during periods without observation. From a systems perspective, this demonstrates that uncertainty is not an issue rooted in sensing limitations, but a structural characteristic of space-based monitoring architectures. Even during ideal conditions where vessel movement is predictable and fulfils steady state assumptions, the lack of consistent observation leads to a rapid buildup in uncertainty parameters. Operationally, as the revisit time increases, the uncertainty grows to a point where predicted trajectories and actions become indistinguishable, leaving the vessel of interest with a wide decision space in which adversarial behaviour can take place. These temporal gaps can be exploited to conduct illicit activity like ship to ship transfer or entry into restricted area, with the possible activities to be conducted growing as revisit time increases, making it one of the most critical parameters for monitoring and deterring hybrid maritime action.

The results also show how improving satellite characteristics like sensor resolution alone does not resolve the system's limitations. The constraint imposed by resolution and coverage being inversely related means that high resolution systems will be characterised by a higher revisit time unless it can be mitigated by the use of constellations, as well as through the integration of data fusion with continuous sensing systems such as AIS or RF based sensors. While current maritime intelligence systems depend increasingly on constellation-based systems with multi orbit configurations, achieving continuous, high resolution global coverage remains challenging, highlighting the need for predictive analytics to be integrated in order to interpolate behaviour between observations and overcome the unavoidable temporal gaps. As shown in the ARVIN case study, vessel tracking between observation windows must be treated as an inference problem instead of a question of measuring positions, shifting the focus of surveillance systems from observation based to predictive tools developed to support decision making through functions like real time confidence scoring for vessel behaviour based on the observed and inferred data.

In addition to the issues outlined above, the rapid development and scaling of hybrid operations is set to reduce the effectiveness of space-based surveillance systems through the offset in Benefit-Cost ratio. Although reusable, modular platforms like the X-37B or reusable rockets are being tested and implemented, the cost of launching satellites into orbit is disproportionately higher compared to the cost of conducting traditional hybrid operations, meaning that the expansion of adversarial vessels is set to outpace the development of surveillance architectures while creating an economic imbalance between the adversary and observing state.

In conclusion, the results show that while extremely useful as an uncooperative surveillance mechanism (producing objective, readily analysable data), surveillance through satellite platforms needs to be sustained by other collection methods and is not viable as a standalone system due to its structural constraints. The uncertainty surrounding the

use of space-based systems warrants for a transition from single source monitoring approaches while integrating probabilistic frameworks capable of interpreting contexts with limited information in complex operational environments like congested chokepoints.

2.3.3. Data Fusion Architectures.

By analysing the limitations identified in satellite-based surveillance and explored with the ARVIN case study, particularly related to the dependency on revisit cycles, environmental constraints or data latency, it can be seen that modern surveillance systems increasingly rely on data fusion architectures to integrate multiple sensing platforms to produce higher quality data analysis. Instead of treating satellite systems as standalone data gathering platforms, as was the case for denied areas during the early cold war satellite programs, these architectures combine cooperative and non-cooperative inputs to transform raw data into a coherent operational database. This standalone methodology was the case for early CORONA missions for example, which operated in isolation and carried single panoramic cameras and returned captured film to Earth via a recovery capsule, later collected in flight by aircraft such as the Lockheed JC-130Bs from the 6594th Test Recovery Control Group. It was not until the introduction of the introduction of the KH-7 GAMBIT satellite series where data fusion could be introduced, with early wide-area search satellites identifying areas of interest, then targeted by high-resolution satellites.

Fig. 6. Recovery of a KH-7 GAMBIT capsule.

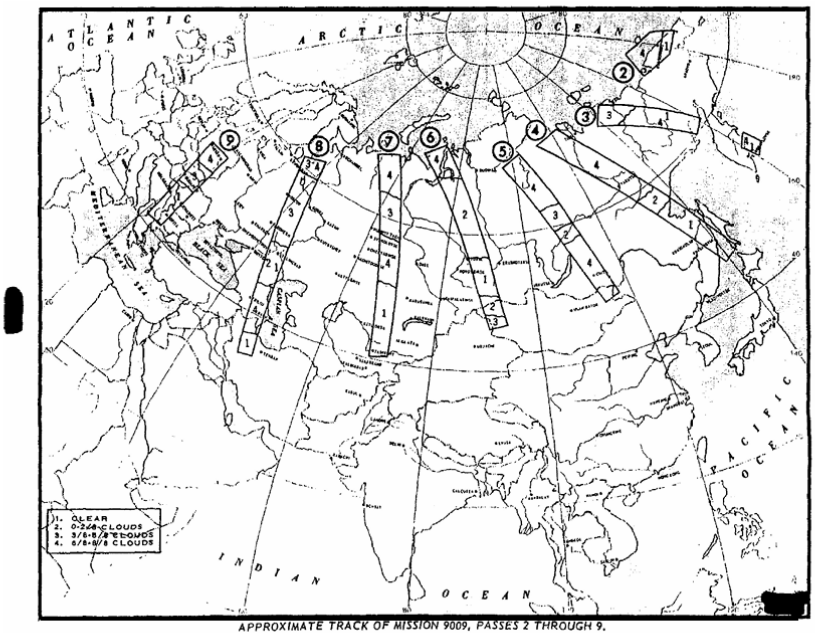

Fig. 7. Mission track for CORONA mission 9009, 18th August 1960. Coverage includes denied areas such as the Kapustin Yar Missile Test Range or the Sarova Nuclear weapons sensor.

From an engineering perspective, this reflects a shift from deterministic single source monitoring towards a probabilistic multi-sensor inference framework, capable of operating under uncertainty and provided with incomplete information. This can only be achieved through the correlation of varied data streams along the three main sensing dimensions, their spatial, spectral and temporal components, enabling the reconstruction of vessel activity in complex environments such as the chokepoints, where single-sensing platforms may provide partial or ambiguous observations. As a result, data fusion not only has become a necessity to complete satellite surveillance systems, but as a necessary evolution in order to achieve reliable situational awareness in an operational context.

2.3.3.1. Architectural Overview and Fusion Levels.

In order to achieve standardisation for data fusion architectures, they are typically structured according to hierarchical models such as the JDL (Joint Director of Laboratories) framework, developed by the US Government in the late 1980's as an effort to reduce inefficiencies in research pertaining to C3I (Command, Control, Communications and Intelligence). In an attempt to integrate Army, Navy and Air Force frameworks the group established several Technology Panels with the main organism being the Technology Panel for C3, composed by the technical directors for the C3 labs at the Naval Ocean Systems Center (NOSC), the Rome Air Force Development Center (NOSC) and the Army's Communications-Electronics Command (CECOM). This framework decomposes the fusion process into multiple levels of abstraction, with the lowest being raw sensor data, processed to extract "measurable features", refined as it escalates through the processing chain, producing outputs on object tracks, situational awareness and ultimately behavioural or threat assessments.

Within maritime applications, this layered approach allows the transformation of multiple sensor outputs into a unified operational dataset. Taking the SAR image interpretation process for example, SAR imagery is first processed to detect potential vessel targets, being fused with data from other platforms such as AIS. As the data escalates through the fusion chain, the relationship levels 0 through 2 extract signals and features, assess the observed entity and the observed situation, with level 3 fusion estimating courses of action, events and impacts of those actions, or Scenario/Outcome Assessment. Finally, at level 4, the model evaluates its own behaviour, assessing its signal/feature parameters and relationships. This hierarchical structure allows for a modular system design, where improvements in individual components can be modified within their sensing or processing layers without redesigning the entire system.

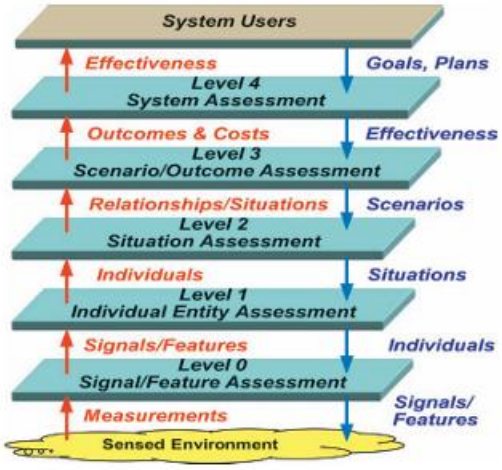

Fig. 8. Proposed 1999 revision for data fusion architectures[29].

[29] Steinberg, A. N., Bowman, C. L., and White, F. E. Revisions to the JDL model. In Proceedings Joint NATO/IRIS Conference, Quebec, Canada, Oct. 1998; reprinted in Sensor Fusion: Architectures, Algorithms and Applications, Proceedings SPIE, Vol. 3719, (1999), 430–441.

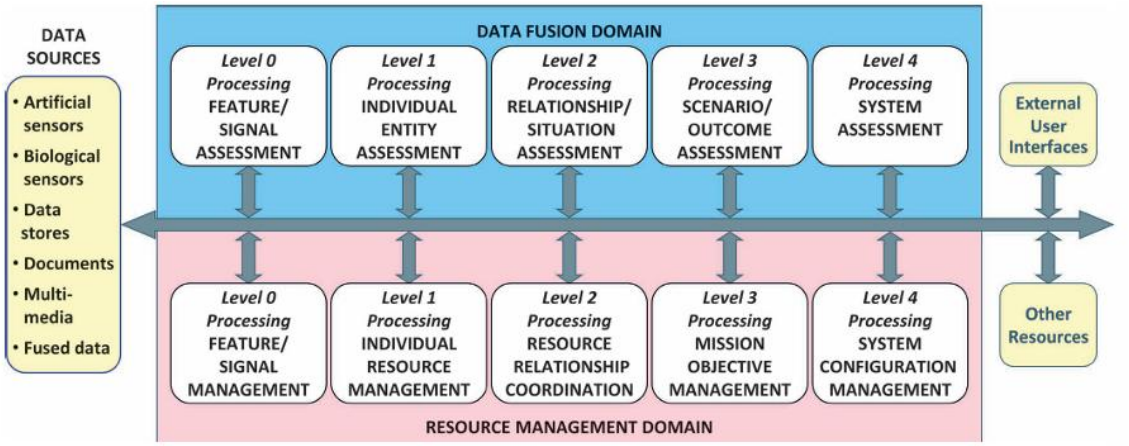

Fig. 9. Data fusion process model[30].

2.3.3.2. Sensor Data Characteristics and Pre-Processing.

One of the main challenges poised by data fusion in maritime intelligence arises from the heterogeneity of sensor data, varying in format, measurement uncertainty, spatial resolution and temporal frequency in accordance with the method by which the information was gathered. For example, while AIS data is received in a structured manner with a determined temporal frequency, SAR or EO systems produce imagery requiring processing and interpretation, with RF sensors producing structured, signal-based physical outputs, but unstructured in a temporal scale, highly dependent on orbital dynamics.

To facilitate the data fusion process, obtained data must be standardised during a pre-processing phase. This process involves alignment in the spatial and temporal dimensions, including georeferencing to align inputs on a common spatial coordinate system and temporal alignment to ensure that the readings correspond to the same time window. This pre-processing phase is also used for signal refinement, applying noise filtering steps or

[30] Steinberg, A., and Snidaro, L. Levels? In Proceedings International Con ference on Information Fusion, Washington, DC, 2015.

feature extraction techniques, highlighting structures from optical systems or relevant emission signatures from RF data for example. Since the fusion process is built upon the pre-processing phase errors introduced during this stage can propagate through the fusion pipeline, making this one of the critical stages for system performance.

2.3.3.3. Data Correlation.

One of the pre-processing steps that is most critical to maritime intelligence applied to chokepoints or critical sea lanes due to the traffic density is the data association process, which determines whether observations from different sensors correspond to the same physical object. One example of this process would be assigning satellite imagery to AIS reported data through pattern recognition and alignment, carried out through a probabilistic framework that attempts to optimise the solution for the discrepancy between predicted vessel states and observed sensor data.

One algorithm designed for such purpose is SUMO (Search for Unidentified Maritime Objects), developed by the European Commission's Joint Research Centre. SUMO accepts imagery from a wide range of satellite SAR platforms, including Sentinel-1, Radarsar-2 or PALSAR-2 among others, making it a versatile tool for maritime surveillance. Once the imagery has been received and pre-processed, SUMA carries our analysis on the individual pixels, assuming that sea noise and clutter follow a K distribution that can be estimated for small tiles. Any pixel with anomalous values is then flagged and grouped together with neighbouring detected pixels, forming clusters to represent ships. The resultant data is then outputted to "Blue Hub" JRC's maritime surveillance research and development platform, which fuses the data received. This is done by overlying reported AIS data on the timestamped

SAR imagery, with a correlation being established between data types if the geographical distance between observations is less than 500 m. This is a necessary filter as imagery is very unlikely to perfectly coincide in the spatial dimension due to shifts in ship speed caused during the SAR image formation.

This process is also frequently complicated by uncertainties in sensor measurements, as well as intentional manipulation like AIS spoofing. As a result, association algorithms must account for process and collection noise as well as adversarial interference, meaning that developing robust algorithms is a significant challenge.

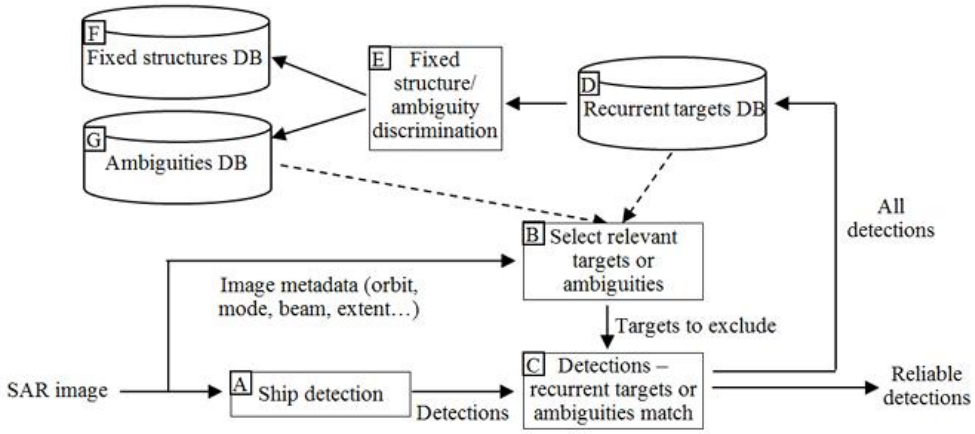

Fig. 10. Block diagram of the recurrent target analysis framework.[31]

2.3.3.4. State Estimation and Tracking.

Once the data stream has been processed and integrated into a single reference framework, the system must estimate the position, velocity and heading to form a dynamic state model for each target. This is achieved through the use of motion models and filtering

[31] Santamaria, C., Alvarez, M., Greidanus, H., Syrris, V., Soille, P., & Argentieri, P. (2017). Mass Processing of Sentinel-1 Images for Maritime Surveillance. *Remote Sensing*, *9*, 678. https://doi.org/10.3390/rs9070678

techniques that estimate the desired parameters from the incoming observations.

One of the main tools for this purpose is the Kalman Filter explored in section 2.3.2.2, useful for estimating systems assumed to be linear. For more complex environments, particle filters can be used to predict behaviour in non-linear environments, where noise is introduced at random not following a gaussian distribution.

Particle filters model shipping behaviour by representing possible positions as a point, resulting in a particle cloud. Noise is added to each particle to account for uncertainty. As a regular Kalman filter, the model goes through an update phase once new sensor readings arrive, with the algorithm comparing it to each particle "forecasting". After the update phase, the particles are given a weight depending on the accuracy of the prediction, with stronger particles being cloned and weaker particles being deleted. This allows the system to model the uncertainty around the vessel movement, resulting in a dense prediction cloud after some iterations. Applied to hybrid operations, readings deviating from the point cloud (contrasted and understood to be real during the pre-processing and assimilation phase) can be treated as anomalous, flagging the ship as suspect. Another advantage of these filters is the resulting smoothed trajectory even in the presence of missing or noisy observations.

2.3.3.5. Anomaly Detection and Behavioural Modelling.

In order to apply data fusion to anomaly detection, a baseline behavioural model needs to be established, allowing for the comparison of the extracted data. These baselines can be derived from fusing a collection of historical data, including typical shipping routes, historical AIS data and operational patterns. This data ledger must also be updated regularly to avoid falling

behind with regards to tactics and standard operating procedures, avoiding problems downstream.

From a data analysis perspective, anomaly detection involves extracting the main behavioural characteristics from vessel trajectories, including parameters such as speed variation, heading changes deviation from established routes. Contrasting algorithms range from more simple outlier detection methodologies, including clustering algorithms to group similar behaviours, to more complex machine learning model capable of understanding vessel behaviour accounting for temporal correlation and dependencies. These algorithms support one of the key advantages of multi-sensor fusion, detection across individual data streams instead of along them. For example, inconsistencies between AIS reported positions and SAR imagery detection may be an indicator of spoofing, while the absence of AIS and RF emissions might indicate deliberate concealment.

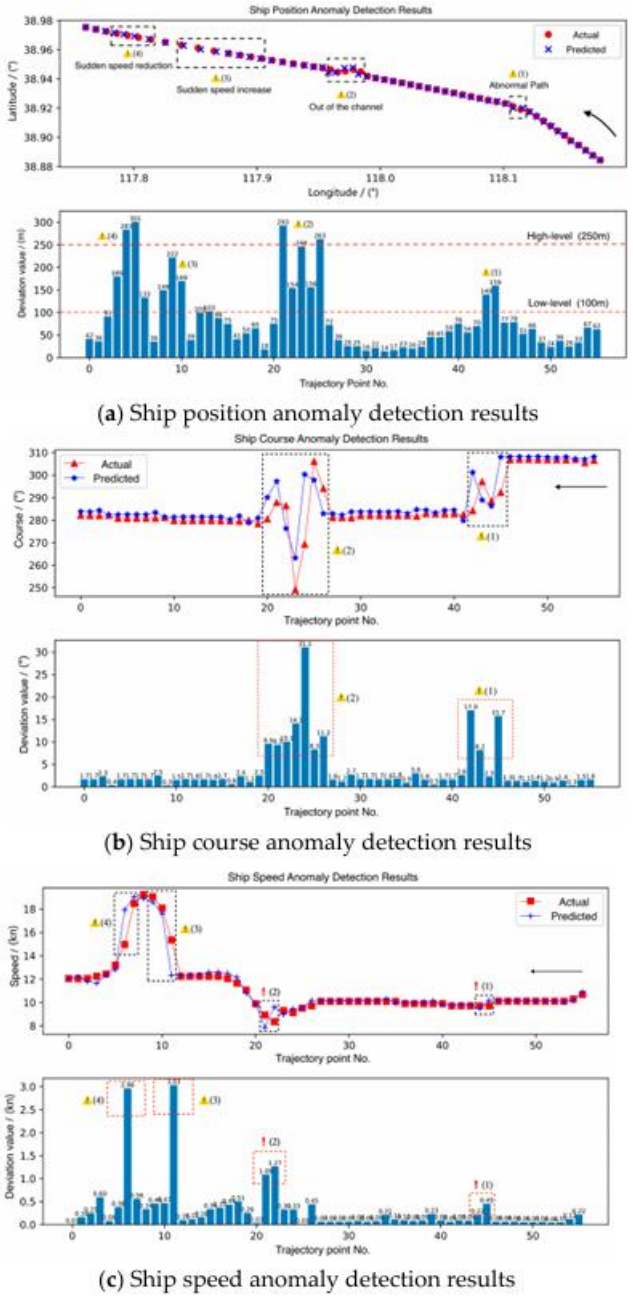

Fig. 11. Example of vessel anomaly detection for position, course and speed dimensions[32].

[32] Zhang, B., Hirayama, K., Ren, H., Wang, D., & Li, H. (2023). Ship Anomalous Behavior Detection Using Clustering and Deep Recurrent Neural Network. *Journal of Marine Science and Engineering*, *11*(4), 763. https://doi.org/10.3390/jmse11040763

2.3.3.6. Consistency Across Multiple Domains.

One of the most significant advantages of utilising data fusion for data examination is the ability to ensure consistency between data sources as explored in section 2.3.3.5, applying pre-processing techniques to ensure a baseline is established among all domains. This allows detection of any vessel state readings that are out of line with the others, including mismatches in vessel position, state and navigational characteristics extracted from sources such as AIS, SAR, EO/IR or RF sensors.

This can be understood as a way to offset the cost-benefit ratio (CBR) of hybrid operations, as maintaining deception requires a much more sophisticated manipulation of data across multiple independent systems, raising the access barrier to maritime adversarial behaviour. It also significantly reduces plausible deniability, as data inconsistencies are a clear indicator of manipulation. As a result, data fusion does not work to eliminate threats entirely, but is highly effective at increasing the cost and complexity of deception, being one of the main mechanisms to shift operational balance.

2.3.3.7. System Constraints and Role in Moden Maritime Intelligence.

Despite the capacity of data fusion to enhance maritime intelligence analysis processes significantly, the data fusion architectures are constrained by some engineering challenges such as the need of large computational resources. The integration of high volumes of data across multiple data sources requires significant amounts of compute, especially when real time processing and threat analysis are required. As a result, intelligence agencies have been forced to outsource their computing capabilities, with the CIA accessing cloud computing through their C2E contract with AWS, Google and IBM among others, amounting to upwards of 10 billion USD, or the NSA developing their Hybrid Compute initiative,

with a 10 billion USD contract being awarded to AWS[33]. This in turn leads to other challenges, such as facing post-quantum cryptography. By moving away from traditional, physical servers, intelligence agencies become more exposed to cyber threats, with billions of dollars being invested in developing these technologies by organizations such as In-Q-Tel, the CIA's non-profit venture capital firm. With the price of technological development increasing as current compute limits are hit, the offset in cost-benefit ratio introduced by the technology is once again offset. Additionally, scalability will add to this constraint, requiring processing millions of datapoints per day across different operating theatres.

To address these challenges, modern systems will require prioritising high-risk targets detected through preliminary analysis. These constraints highlight the need to develop efficient algorithms and streamline interpretation processes to ensure data fusion architectures remain viable at the required scale. However, data fusion systems represent the core of the analysis systems allowing for effective data interpretation, making their development non-negotiable. By integrating multiple data points from a variety of sensing streams, they create an integrated operational picture, providing the clarity necessary to take informed decisions at a strategic level, unachievable by analysing potentially ambiguous information obtained from single sensing sources. Ultimately, the effectiveness of maritime intelligence systems is not exclusively determined by the quality of the sensing platforms, but by the effectiveness of the fusion architectures integrating them.

2.3.4. Machine Learning and Behavioural Modelling.

After the integration of data streams to form a coherent, heterogeneous packet through data fusion, the next stage in maritime intelligence systems involves extracting data patterns and inferring behaviour. This is achieved through the application of machine learning techniques, to establish baseline behaviour and

[33] NSA re-awards $10B WildandStormy cloud computing contract to AWS | FedScoop

detect deviations from it. It extends beyond spatial correlation, detected in the pre-processing and integration phase, introducing behavioural modelling, where vessel intent is predicted through the temporal analysis of state vectors.

2.3.4.1. Feature Extraction.

In order to carry out the analysis, data from the combined sensing platforms is classified into different vectors depending on their nature. As an example, these can be kinematic (including parameters such as Speed Over Ground (SOG), Course Over Ground (COG) or Rate of Turn (RoT)), positional (including the four features describing a vessel's position as described by Laxhammar (2008)[34], momentary position in latitudinal and longitudinal coordinates and the velocity in latitudinal and longitudinal direction), or Contextual, storing the vessel's state in context allowing for the detection of abnormal situations, such as abnormal routes in a temporal context (related to seasons, days of the week, workdays vs holidays…)[35], or vessels taking unusual routes. Maritime operation analysis can store further data such as emission based data capturing RF activity or AIS detected anomalies.

This stage defines the features that learning algorithms will operate on, making it a critical step in the analysis chain. Poor choices of features can degrade model performance and hinder the correct operation of the system, while organised and clearly defined feature sets can enable accurate predictions even in noisy settings or incomplete datasets.

2.3.4.2. Probabilistic Trajectory Modelling.

While classic filtering techniques such as the Kalman filters explored are effective for linear state predictions, they are less accurate when

[34] Laxhammar, R. (2008). *Anomaly detection for sea surveillance*. 1–8.

[35] Hayes, M. A., & Capretz, M. A. (2015). Contextual anomaly detection framework for big sensor data. *Journal of Big Data*, *2*(1), 2. https://doi.org/10.1186/s40537-014-0011-y

examining non-linear temporal dependencies. This is the area where sequence based models excel. Two models that have been included in previous prediction software include Convolutional Neural Networks (CN) and Long Short-Term Memory (LSTM) networks, with both of them having slight differences that make them more valuable for particular applications.

While CNN models are good at analysing short term patterns, LSTM performs better when analysing long term dependencies and evolving trends over time, having long-term memory capabilities. However, LSTM require large datasets for effective training and are developed through a slower training. Despite these advantages, LSTM networks are considered the more suitable type of model for predictive analysis, as it enables the detection of recurring operational patterns, such as speed reductions in congested chokepoints or course changes in established routes. Anomaly detection is then performed by comparing predicted trajectories with observed states, with deviations being quantified to provide a probabilistic indication of abnormal behaviour.

2.3.4.3. Behavioural Classification and Intent Inference.

Once anomalous behaviour has been detected, the fused data can be analysed using classification models such as Random Forest or Support Vector Machine (SVM) algorithms, two of the most widely used in maritime intelligence. These models associate anomalous patterns in the data to specific types of behaviour, allowing the system to move from highlighting anomalous behaviour to understanding their potential meaning.

A Random Forest algorithm is composed of decision trees that split the received data based on feature thresholds (for example, if the distance between vessels or the speed difference between vessels is lower or equal to a certain value), with each branch ending in a classification (e.g. Normal behaviour or loitering). The forest works by combining the output from

the trees and outputting the majority vote, which can have a confidence score attached to it depending on the decision tree classification split. Some advantages of this algorithm when being applied to maritime intelligence is it works well with tabular data, it is easy to interpret and withholds its integrity when facing incomplete or noisy data, however, since its inputs are static feature representations and are not time based, it does not capture temporal dependencies. This means that time-based behaviour must be encoded manually into the model using lags or timestamps.

Support Vector Machine algorithms work by compiling past examples of labelled collected data and establishing the boundary that differentiates normal and threatening behaviour. It does this by creating a hyperplane (a flat (n-1) dimensional subspace that divides a n-dimensional space in two) defined by support vectors, located using only the most important points. It works well with high-dimensional spaces, an advantage for environments with complex dynamics such as the grey zone, but does not scale well with larger datasets and, like Random Forest algorithms, is not inherently temporal, as well as being harder to tune.

2.3.4.4. System Role and Implications.

The classification layer is what allows the system to move from anomaly detection to supporting intelligence on an operational level, interpreting threats and intentions. By integrating machine learning into the data fusion pipeline, support systems evolve from passive monitoring tools, identifying patterns, detecting deviations and generating risk indicators instead of solely relying on human analysts to interpret large volumes of data.

This is particularly significant in this case since adversarial action is designed to remain ambiguous when seen through a single data source, with complications arising when attempting to fully conceal activity from

every different data source. In this regard, behavioural modelling does not eliminate uncertainty entirely but quantifies it within the model's boundaries, aggregating behavioural deviations across the spatial, temporal and spectral dimensions to support more confident attribution.

2.3.5. Open-Source Intelligence (OSINT).

While machine learning models extract behavioural patterns, it is important to add contextual information for an accurate interpretation within a broader framework. Open-Source Intelligence (OSINT) provides context by integrating data mainly aggregated from the internet, enabling attribution and supporting identification by enabling access to a variety of human (HUMINT) or technical sources.

2.3.5.1. Contextual Data Integration.

Maritime OSINT is mainly characterised by the use of existing datasets to complement information obtained from sensing capabilities. These datasets include information registries containing vessel identifiers (MMSI, IMO), or historical ownership and management records. Additionally, vessel regulatory compliance status can be verified through the integration of sanctions lists (e.g. OFAC or EU dynamic lists).

This can prove to be challenging due to datasets being presented in inconsistent data formats or ambiguity across sources. Manipulation of vessel identifiers such as MMSI spoofing or organisational concealment, for example through the use of shell companies, can further complicate the establishment of a reliable reference framework. As a result, OSINT systems must be first pre-processed and filtered to ensure that fused records refer to the same vessels, forming a consistent baseline.

2.3.5.2. Automated Information Processing.

Due to the scale and heterogeneous presentation of OSINT data, modern maritime intelligence systems rely on automated processing pipelines based on Extract, Transform and Load (ETL) architectures. Through this methodology, raw data is progressively refined into consistent, standardised data that is useful for analysis.

This process also enables the expansion of the data dimensions, integrating non-kinematic features such as ownership structures, historical behavioural indicators or previous AIS anomalies. These systems also manage to mitigate the information overload problem by transforming dense inputs into manageable and analysis ready data streams, highlighting concerning data rather than importing unusable data that would make the decision system slower.

2.3.5.3.Technical and Operational Attribution.

The primary role of OSINT in maritime intelligence is to support attribution in a highly complex environment, which it can do in two distinct ways.

In first place, OSINT helps by supporting technical attribution, identifying vessels based on verifiable data such as registry information. This allows for an immediate flagging of suspected threats, saving compute. On the other hand, it also supports operational attribution by associating recorded historical activity to extracted behavioural patterns to infer relationships and intent. By linking anomalous behaviour to external datasets, OSINT provides the basis required to move from detection to interpretation. Like the other technological systems explored in this section, it does not eliminate uncertainty entirely, but it reduces the ambiguity of the environment, supporting robust intelligence assessments and increasing the cost of grey zone activity.

2.3.6. Graph-Based Intelligence and Neural Networks.

Graph-based intelligence introduces a structural layer to maritime intelligence, enabling the analysis of relationships between vessels and operating entities, supporting the identification of ownership and attribution. Instead of treating vessels as isolated units, this approach models the maritime domain as an interconnected network, allowing the detection of coordinated activity and hidden dependencies.

2.3.6.1. Modelling of Networks.

By representing maritime activity as a network composed of nodes (V, corresponding to entities like vessels, ports or infrastructure) and edges (E, representing the relationships between them such as repeated interactions, ownership or repeated coincidence in location) the maritime environment can be represented as a structure instead of individual entities.

This representation enables understanding structured movement patterns across the observed dimensions, allowing for a more complete understanding of maritime activity, particularly in dense environments such as chokepoints, where interactions between vessels play a critical role in forecasting.

2.3.6.2. Representation Learning and Graph Neural Networks.

To analyse the relationships, Graph Neural Networks (GNN) are used to extend machine learning techniques to graph-based data. These models generate low dimensional representations of nodes and networks by incorporating both their individual features and their connections, modelling the interactions of vessels with neighbouring entities. This is done through the application of architectures such as GraphSAGE (Graph Sample and aggreGatE), especially useful for this application as it is tailored to understand relationships in a network that is constantly changing and expanding.

In order to analyse the network, GraphSAGE samples the immediate network around the examined ship (e.g. recently visited ports or ships in its proximity) rather than the whole network, allowing it to function in a data-dense environment. Following the sampling, it derives a function to average the behaviour of the vessel's neighbours, updating the vessel's profile based on its surrounding nodes to predict if it is a threat. As a result, graph-based learning captures local interactions such as vessel proximity, allowing it to detect behaviour such as shadowing or swarming[36], with broader patterns, improving the ability of the system to detect coordinated behaviour.

2.3.6.3. Graph-Based Anomaly Detection.

Graph-based architectures can also be applied to anomaly detection, by focusing on identifying deviations from expected network structures and interaction patterns. Much like other machine learning algorithms, the deviation of nodes from current network states is derived from a learned baseline, identifying irregularities through statistical approaches.

These techniques are particularly effective in detecting coordinated activities such as ship-to-ship transfers, which may appear normal when analysed at the individual level. However, graph-based detection allows analysts to have a clear vision of networks at a higher level, revealing network anomalies and highlighting relationships that would otherwise be hidden in the high-volume data streams, a feature hybrid operations rely on by definition.

2.3.7. Predictive Risk Scoring Systems.

As data flows through the processing pipeline, it transitions from data processing to interpretation, allowing for proactive threat monitoring. One

[36] Signals in the Swarm: The Data Behind China's Maritime Gray Zone Campaign Near Taiwan

form of data interpretation that is especially interesting is predictive risk scoring. Instead of analysing isolated observations, these systems assign a dynamic risk score to each vessel by quantifying deviations from established behavioural baselines. This not only allows data to be displayed in a visual format that is easily interpreted but also allows for a fast fine-tuning of parameters and weights according to the scenario or as tactics and threats evolve.

2.3.7.1. Aggregation from Multiple Sources.

The basis of predictive risk scoring lies in the aggregation of data at the end of the intelligence pipeline, after being processed. These inputs include data from cooperative reporting systems such as AIS, Long Range Identification and Tracking (LRIT) or Vessel Monitoring System (VMS) as well as non-cooperative sensing platforms such as SAR, radar or RF detection as well as contextual datasets. One advantage of the integration of multiple datasets in these systems is the ability to fine tune weights by data type depending on the circumstances, lowering the importance of self-reported data if it is discovered to be misleading for example.

2.3.7.2. Probabilistic Risk Modelling.

Even with a robust data pipeline, risk assessments for prediction systems are inherently probabilistic, and are heavily influenced by uncertainties created by incomplete observations, sensor noise or adversarial behaviour. Therefore, predictive risk scoring systems rely on machine learning models to estimate the likelihood of vessels engaging in threatening behaviour.

Some approaches include density-based models such as Gaussian Mixture Models (GMM) or Kernel Density Estimation. Density-based approaches operate on the assumption that normal activity creates clustered patterns, with irregular being identified as stand out profiles in low density areas, with threat scores being assigned by measuring the deviation of a vessels

state from the normal baseline, with a higher deviation getting assigned a higher risk score. While density models highlight vessels that could pose a threat and quantify the risk, the application of Bayesian frameworks can make inference possible, deriving potential threat states from behavioural data and contextual indicators. This is done by the application of Bayes' Theorem shown as follows.

$$P(Threat|Evidence) = \left(\frac{P(Evidence|Threat) \times P(Threat)}{P(Evidence)}\right)$$

Or translated into a problem statement, the probability of the examined vessel being a threat given the evidence at hand depends on the how likely a threatening vessel would show the evidence at hand (spoofing or AIS disruption are telltale signs), the probability of the ship being a threat derived from prior data (ship's flag, ownership network…) and how common the evidence is in comparison to the rest of the fleet.

State estimation techniques such as Kalman Filters or Particle Filters can also be used for risk modelling as seen in previous sections, incorporating process noise to account for uncertainty in the system, with deviations between predicted and observed states being transformed into a risk estimate.

2.3.7.3. Dynamic Risk Scoring and Temporal Updating.

Given the dynamic nature of the maritime environment, assigned risk values must be continuously updated as new observations become available. Therefore, predictive risk scoring systems must be recursive in nature.

In order to carry out these updates, sliding time windows (where only a fixed duration of past behaviour is looked at, with older data points being dropped as newer ones get introduced into the system) or Sequential Updating Schemes, where risk before new evidence being introduced is

used as a base for the new risk calculation (one example of a sequential updating algorithm is the Kalman filter). Additionally, as new inputs are integrated, they are given higher weights than the older data points to avoid normal historical data skewing the system into not recognising sudden threatening behaviour. For example, a vessel following a normal route may initially be assigned a low-risk score, but abrupt course variations or loitering can rapidly increase its risk score.

2.3.7.4. Decision Thresholding and Operational Integration.

Since the effectiveness of outputs at the end of the intelligence chain is more closely related to their integration into operational assessments, it is imperative that risk estimation tools are clearly defined through the use of decision thresholds.

These thresholds are a way of calibrating sensitivity to protect the system against false alarms among other functionalities. Multiple threshold levels can be set as a way to clarify response types. While low level alerts would signal the need to document the incident and contribute to building passive deterrence by identifying suspicious patters, high level alerts can be used to indicate the gathering of enough conclusive evidence to support escalation and a coordinated response. Ultimately, the integration must be treated as an operational aid, with political attribution remaining a human policy decision, with technology supporting rather than replacing human judgement.

2.3.7.5. Strategic Role in Hybrid Maritime Operations.

As seen in this section, predictive risk scoring systems are critical to challenging plausible deniability and reducing ambiguity. This is done by providing an overview of vessel characteristics, aggregating inconsistencies across behavioural and contextual-derived data. From a systems perspective, deception under this kind of surveillance would imply maintaining consistency across all of the monitored sensor streams,

across the spatial, spectral and temporal dimensions, increasing the cost of action and significantly complicating grey zone operations.

2.3.8. Cloud Computing and Real-Time Processing.

The scale of modern maritime operations, combined with the computational demand of processing the large volumes of sensing data demand a transition from localised computing infrastructures to distributed cloud-based architectures. Only through the use of these systems can the necessary scalability to receive, store and process the dense data streams while supporting real time analysis.

By using cloud computing, maritime intelligence moves from static hardware to flexible environments where computing power can be scaled up or down as desired. This is especially important from a cost-benefit analysis perspective, as no overcapacity has to be built into the system, wasting resources during low-demand periods. Instead, agencies can focus on maximising computational resources and optimising its use between intelligence domains, often within the same agency.

2.3.8.1. Edge Computing and Distributed Processing.

Given the scale of maritime hybrid operations, which have now expanded into a global theatre of operations, one of the main interests of analysis systems is to perform initial data analysis closer to the data source in order to reduce data processing delays.

This is achieved through “edge computing”, a strategy in which the initial interpretation and data formatting (including tasks such as signal conditioning, feature extraction and preliminary anomaly detection) are carried out in local processing centres, aboard collection vessels, unmanned platforms or coastal stations. This distributed processing ensures that only relevant information is passed along down the data

stream, reducing bandwidth requirements and facilitating scaling systems as operations require it. Fog computing, or intermediate layers, can also be introduced to further streamline the process, taking in data at a regional rather than local scale. This ensures a flexible, responsive procedure that is fit to respond to fast paced modern hybrid operations.

2.3.8.2. System Constraints.

Despite their advantages, cloud-based maritime intelligence systems are subject to several constraints, such as bandwidth limitations. The cost and limited availability of satellite communications data collection and interpretation architectures rely on means that trade-offs between data resolution and latency must be reached. Furthermore, cloud infrastructures also introduce cybersecurity risks, as they become potential targets for hacking attacks or communication intercepts. This risk will be further accentuated by the advent of post-quantum cryptography, making the development of secure cloud environments a priority for government agencies across the world. However, it is important to mention that distributed cloud infrastructures also diversify risk by not hosting systems in a single physical location, providing more security against kinetic operations (such as sabotage or destruction of physical servers), albeit at an increased vulnerability to cyber/electronic warfare.

Additionally, the reliance on access to large scale computational resources reinforces existing challenges related to scalability, or energy consumption. As data volumes continue to grow, it is imperative that processing pipelines are optimised and governed by efficient resource allocation to maintain system performance, with guaranteed access to secure energy supplies.

2.3.9. Digital Twins and Simulation Environments.

Digital twin technology, pioneered by NASA in the early 1960's to mirror spacecraft behaviour during the Apollo missions[37], was only fully developed and implemented around 2010, when NASA developed digital twin models or physical assets. These systems enable simulation, prediction and analysis in a secure manner, without direct interaction with the physical world.

Unlike traditional modelling environments, digital twins are continuously evolving, synchronising with the real world in a way that has only been enabled due to developments in sensor integration and data fusion. In the maritime domain, this allows virtual representations of vessels, shipping lanes or operational theatres to evolve as data is received from updated satellite imagery, AIS transmissions or contextual data is received.

Instead of solely displaying observations, modern systems must attempt to construct predictive representations of reality, capable of estimating future states and identifying vulnerabilities by evaluating scenarios before they physically occur. This is particularly relevant in the context of this report, where adversarial behaviour is specifically designed to remain ambiguous and difficult to attribute when observed through a reactive lens.

This evolution also reflects a wider conceptual shift in modern intelligence systems towards simulation environments or world models, where the objective is no longer to observe the operational space but to build live digital environments through forecasting and interpolation. As explored by Ed Parsons in "The Map of Dreams : Why AI's "World Models" might be the Geospatial Industry's Ultimate Disruption"[38] advances in AI-driven world models are shifting geospatial systems from static cartographic representations towards predictive environments. Instead of representing the observed reality through maps, these systems reconstruct incomplete environments from fragmented databases, which in the case of

[37] https://web.mit.edu/digitalapollo/Documents/Chapter7/fellemanhybrid.pdf
[38] The Map of Dreams : Why AI's "World Models" might be the Geospatial Industry's Ultimate Disruption – edparsons.com

maritime surveillance, could enable the transition from retrospective analysis towards probabilistic estimation of vessel behaviour.

2.3.9.1. Predictive Analysis and Operational Integration.

One of the main advantages of digital twins is their ability to support simulation-based analysis by generating multiple possible scenarios. In this manner, these systems allow intelligence analysts to explore potential outcomes under a variety of conditions and assumptions.

This is done through the use of simulation techniques such as trajectory propagation or stochastic sampling (for example, using Monte Carlo simulations), estimating the range of system states, or possible outcomes, under uncertainty. Through repeated simulation, digital twins can identify potential risks such as congestion, anomalous vessel behaviour or infrastructure vulnerability.

In intelligence context, digital twin architectures are used to support operational monitoring and strategic analysis by integrating real time data with predictive models. They are particularly useful in environments where labelled data is scarce, with virtual environments being used to train and validate machine learning models. In this manner, systems can improve their detection capabilities and precision without relying on real life events for training.

Additionally, digital twins can be used to support maintenance and optimise existing systems by identifying deviations from expected system performance, functioning as a tool for process improvement.

2.3.9.2. Simulacra and Simulation.

As machine learning and predictive architectures become increasingly central to analysis, an important distinction between observing operational

reality and simulated representations of it must be made. This issue has been discussed in depth in a number of papers, most notably explored by French philosopher Jean Baudrillard in Simulacra and Simulation, in which he examines the relationship between real and derived environments, and how representations of reality are capable of evolving to replace or obscure the physical world.

This creates an interesting dilemma for the use of digital twins as, by relying excessively on simulated environments with a lack of real world grounding, operational decisions may be detached from reality. This occurs when the analytical authority of the environments is inverted, with simulation results being taken as absolute truths instead of probabilistic inferences. The risk becomes more significant when the data connection between environments becomes heavily weighted in the reverse direction, with more data flowing from the simulated world to the real scenario. When the data connection element becomes a closed feedback loop in which the assumptions generated by predictive algorithms are executed upon without a validating effort, the simulation stops being a tool for understanding reality and gradually replaces it in the analyst's perception.

In this sense, the simulation becomes a simulacrum, an environment in which the system infers reality from its own expectations rather than the examined theatre of operations. This concern extends beyond technical inaccuracies and introduces a conceptual debate at the foundation of modern intelligence interpretation systems, where increasingly autonomous systems are becoming capable of influencing strategic decisions without human verification.

2.3.9.3. System Constraints.

Overall, the effectiveness of digital twin systems is constrained by the availability and quality of input data. In regions with limited sensor coverage or unreliable data streams, the accuracy of the model may degrade, leading to increased uncertainty.

There is also a trade off between model accuracy and computational cost, with high resolution requiring large amounts of computing power, limiting their application to large agencies and increasing the barrier of entry. Finally, while digital twins provide analytical support, it is important to consider that human interpretation is invaluable, and the attribution of intent or responsibility still requires human judgement, especially in contexts where operations are designed to be ambiguous and can be designed to exploit model inaccuracies.

2.3.10. Emerging Intelligence Platforms.

As hybrid operations evolve beyond strategic disruption into more complicated, decentralised and ambiguous campaigns, intelligence systems are developing towards alternative collection and analysis methodologies. These systems expand the intelligence cycle outside conventional frameworks to access new data pools, and in doing so, increase the cost of conducting grey zone operations, which must remain ambiguous among all domains to remain covert.

This transition has been favoured by the data fusion environment, where fragmented indicators are interpreted to form a complete operational picture. As a result of this evolution in intelligence theory, extracting information from seemingly unrelated datasets has become a research priority for intelligence agencies.

2.3.10.1. Ubiquitous Technical Surveillance (UTS).

One of the defining features of the modern information environment is the sensor saturated environment, with exploitable intelligence data being generated by connected devices in our surroundings. This is known as Ubiquitous Technical Surveillance and comprises the technical collection of data obtained from digital infrastructure such as smartphones, wearable

devices, IoT devices or even vehicle telemetry. Classified as one of the main challenges faced by intelligence agencies, escaping surveillance is extremely complicated due to Pattern of Life analysis, which classifies behaviour on a “normality spectrum”. This ensures that full anonymity is flagged as an outlier, forcing both agencies and actors to rethink their approaches to tradecraft. In maritime operations, this signifies that vessels attempting to go dark will be highlighted sooner than those acting under apparent normality.

This environment emphasises the need for careful signature management, where even brief lapses in discipline or breaches in operational security (OPSEC) can lead to detection. Well known examples of OPSEC breaches include those derived from STRAVA leaks, a popular app for recording workouts such as running or cycling, tracking metrics including position. In 2018, global heatmaps where made publicly available as a way of showcasing exercise patterns, inadvertently highlighting movement patterns of military personnel operating inside sensitive areas. A more recent example includes the tracking of the French aircraft carrier Charles de Gaulle strike group, demonstrating how data obtained from publicly available systems can have an important operational impact through metadata aggregation.

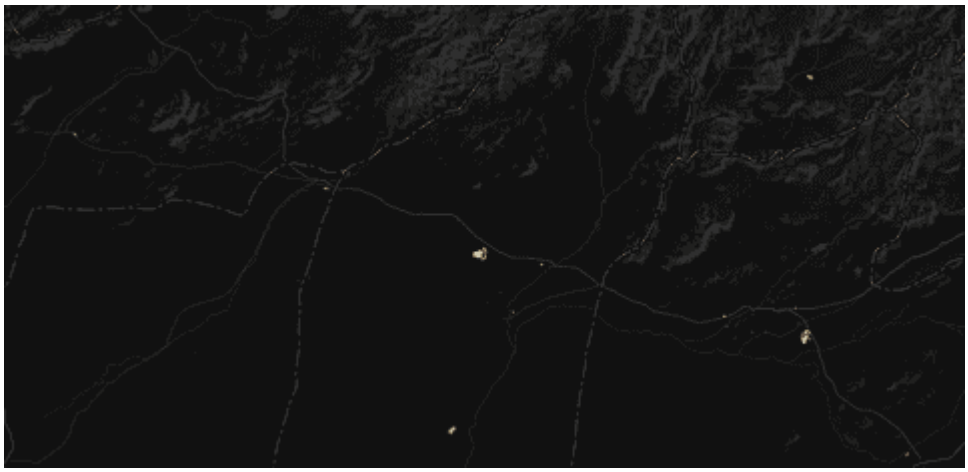

Fig. 12. Military bases in Helmand province, Afghanistan highlighted by STRAVA leaks.[39]

Fig. 13. Charles de Gaulle carrier strike group location highlighted by STRAVA.

This introduces another significant shift in approaches to intelligence, with systems monitoring persistent civilian generated data streams to identify anomalies or correlations instead of tasking sensors towards targets. In turn, this strategy requires scalable processing systems, capable of

[39] Fitness tracking app Strava gives away location of secret US army bases | GPS | The Guardian

analysing heterogeneous data and behavioural signatures at massive scales. In addition, using data logged by civilian systems introduces renewed debates about privacy concerns and the erosion of the boundary between private and public data. Apart from the much explored regulatory contradictions and collateral data collection, referring to the processing of personal data of non-related individuals by surveillance systems, creating permanent records or activity ledgers of third parties, this introduces the question of normalised surveillance. The use of ubiquitous systems risks a societal shift where constant monitoring is no longer seen as an extraordinary measure but a requirement for upholding security, often accepted as a consequence of using online services.

2.3.10.2. Advertising Intelligence (ADINT).

Another relevant, involuntary data collection system that has been the subject of recent studies is Advertising Intelligence, or ADINT. This refers to the large-scale extraction and aggregation of metadata obtained through digital advertising. Unlike traditional intelligence collection systems, ADINT does not rely on sensing infrastructure but instead relies on the continuous flow of information generated by online devices.

Modern advertising systems operate through auction-based systems where digital advertisement impressions are bought and sold in milliseconds as a webpage loads. This allows for targeted advertisement based on user profiles, with the higher bidder winning the right to display an advertisement on the target's device. In order to inform the bidder of the user profile, allowing them to price their bid, information about the targeted individual is released to the advertisers, including geolocation, a unique identifier known as the Mobile Advertising ID (MAID), or other personally identifying information. From an operational perspective, this can reveal behavioural patterns and other data that may reveal incoherences in a well-constructed façade for a hybrid operation. One of the reasons why ADINT has been a subject of increased interest is also due to the accessibility of advertising data, further highlighting the challenges regarding data privacy and reliability.

2.3.10.3. Betting Intelligence (BETINT) and Prediction Markets.

BETINT, or Betting Intelligence, is another emerging intelligence collection system spurred by the rise in popularity of prediction markets, using financial data and predictive contracts to forecast events. Originally proposed by Defence Advanced Research Projects Agency (DARPA) in 2003 as a way to explore awareness in the information age, it was heavily criticised and was cancelled within a few weeks of its public announcement, however, it set the basis for modern prediction market analysis. Prediction markets as imagined by DARPA functioned much like future markets operating in regular financial environments, however, while futures derive their value from the price of an underlying commodity, creating a contract for its future value, prediction derivatives gain value from the likelihood of a future event occurring. In a prediction market, users buy and sell contracts tied to the outcome of current affairs, with the market price being set from an estimate of the event occurring. Once the deadline for the contract has passed, the market is settled by paying out depending on the binary outcome.

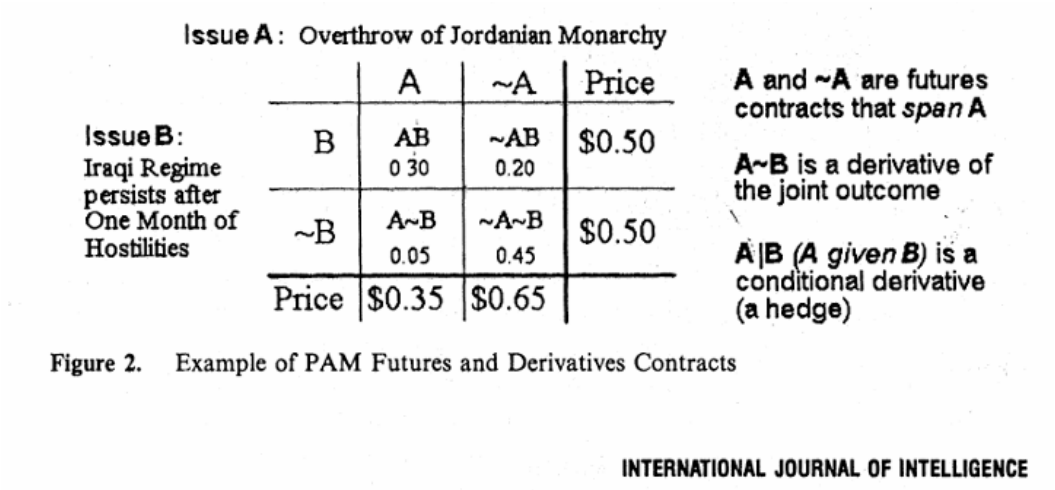

Issue A: Overthrow of Jordanian Monarchy

Issue B: Iraqi Regime persists after One Month of Hostilities	A	~A	Price
B	AB 0.30	~AB 0.20	$0.50
~B	A~B 0.05	~A~B 0.45	$0.50
Price	$0.35	$0.65	

A and ~A are futures contracts that *span* A

A~B is a derivative of the joint outcome

A|B (*A given B*) is a conditional derivative (a hedge)

Figure 2. Example of PAM Futures and Derivatives Contracts

INTERNATIONAL JOURNAL OF INTELLIGENCE

Fig. 14. Example of a PAM Future and Derivative contract. International Journal of Intelligence, 2004.[40]

[40] DARPA's Policy Analysis Market for Intelligence: Outside the Box or Off the Wall?

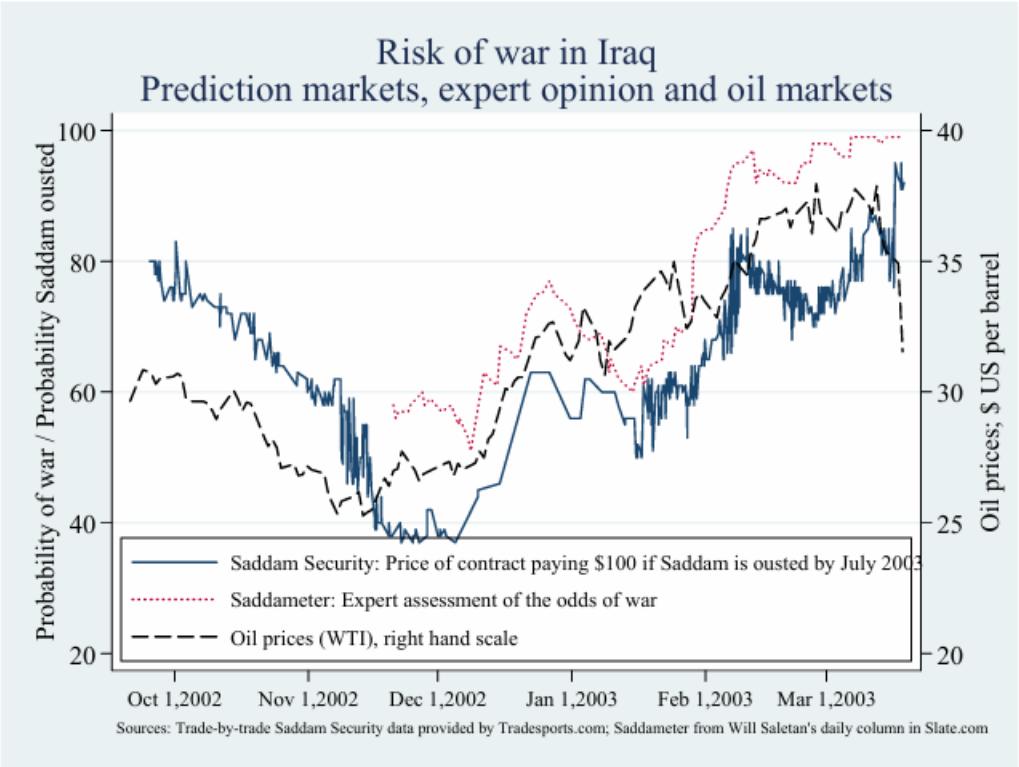

Fig. 15. Commodity vs Prediction Market Future probabilities matching profiles. Prediction Markets in Theory and Practice. Justin Wolfers and Eric Zitzewitz, 2006.

The theory originally backing PAM and modern markets is based on the efficient market theory, according to which securities markets are extremely efficient in reflecting information about individual stocks (outcomes) or the stock market (geopolitical environment). By diversifying the profile of the contributing participants, fragmented information can be aggregated across large populations, contributing to the comprehensive nature of the study. While original plans for MAP were to create a market only accessible to the intelligence community, regulations on inter-agency fund transfer forced it to become open to the public. While this expanded the user base to allow for additional inputs, this caused some concern among the developers. In the first place, an open market meant that adversaries would have access to the information derived from the market, as well as gain insights into US research interests. Another crucial concern was that it introduced the possibility of adversarial manipulation. However, this last concern was quickly dismissed for several reasons, in

first place, it acts as an investment incentive for accurate traders that would drive the market back to equilibrium, resulting in a net loss for the adversary. Furthermore, initiatives like MAP introduced payout caps to ensure that benefiting from financing geopolitical disruptions would be impossible. It is important to note that this theory relies on the assumption that the information flow is equally as efficient and unimpeded, and that future price changes are only dependent on future scenario changes and are independent of present evolution. Therefore, any insider trades and manipulation attempts exploit information asymmetry, causing temporary distortions in contract price that are eventually corrected. However, from a collection standpoint, these manipulation attempts can be framed as highly visible statistical anomalies. This information gathering technique is most relevant with current prediction markets and the emergence of platforms like Polymarket or Kalshi. From an intelligence perspective, the sheer market capitalization provides a valuable data source which can easily be explored in the context of hybrid operations, already represented in open markets.

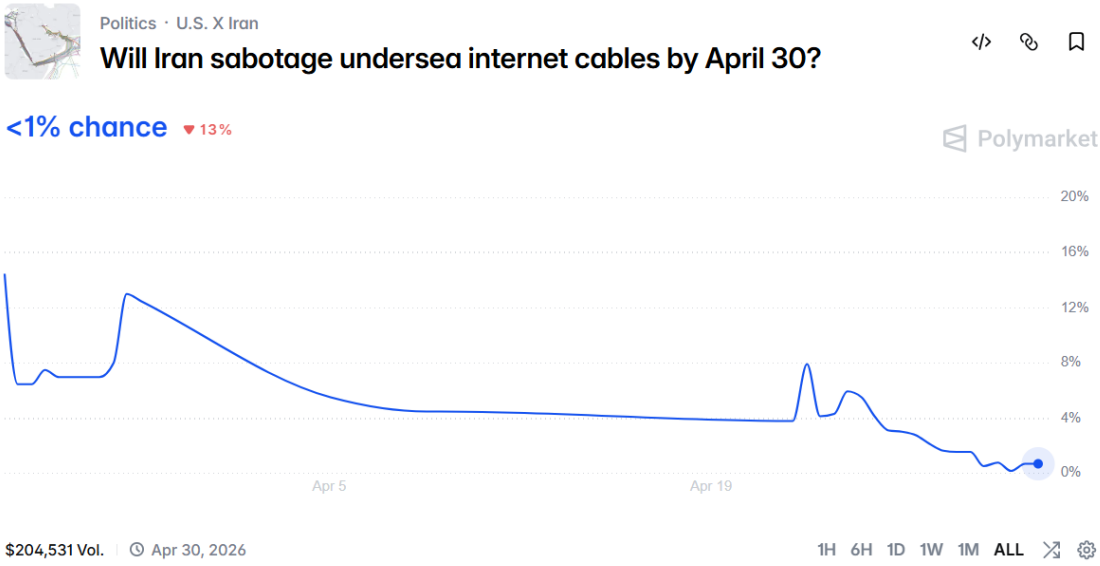

Fig. 16. Prediction Market on undersea cable sabotage on Polymarket, resolved NO with a total market volume of 204531 USD. [41]

[41] Will Iran sabotage undersea internet cables by April 30? ... 2026 | Polymarket

2.3.10.4. Synthetic Anomaly Generation.

While maritime intelligence systems become increasingly dependent on machine learning architectures, one of the main limitations faced by predictive models is the lack of labelled datasets corresponding to hybrid adversarial behaviour. This is due to the nature of hostile activity, which is designed to avoid detection and, although going through a period of expansion, is intentionally rare. This makes model training a challenge that modern intelligence systems attempt to address through the use of synthetic anomaly generation, where artificial behavioural patterns are generated to create and support training datasets, as well as improving the robustness of anomaly detection systems.

By using synthetic generation systems to create simulated trajectories and behavioural signatures characteristic of plausible adversarial conditions datasets avoid depending exclusively on limited historical records. This is especially important in the current state of hybrid operations, where the fast evolution of techniques doesn't allow for comprehensive datasets to be built before becoming outdated. These anomalies can include AIS spoofing, probing, loitering or ship to ship transfers, which are introduced into otherwise normal traffic datasets, allowing machine learning models to train on identifying rare behavioural deviations without waiting for real world events to occur (with the consequences they carry).

The development of this technology is closely linked to the development of digital twin architectures, as they are often integrated into simulation environments, with synthetic behaviour being modelled through probabilistic or AI driven approaches. Recent models have also used Large Language Models into the process, allowing systems to develop situation briefs into operational environments. Synthetic anomaly generation also supports system validation and testing, evaluating the sensitivity of detection software and identifying weaknesses before system deployment. Additionally, this approach also permits the user to fine tune model

parameters and weights to best fit the operational and strategic environment, resulting in a tailored fit for the analyst's convenience.

However, much like digital twins, these approaches also introduce risks related to model bias. If generated anomalies do not reflect the true operational environment, systems may become detached from reality, being tuned to artificial and disconnected behaviour. This reflects the need for synthetic anomaly generation to be understood as a complement to grounded, real datasets and not as a replacement to ensure adequate performance when faced with real world scenarios.

2.4. Geopolitical and Strategic Context of Hybrid Maritime Operations.

2.4.1. The Grey Zone and Maritime Hybrid Warfare.

The maritime domain has become one of the main theatre of operations for grey zone operations in recent years due to the complicated nature of maritime policing. Shared and often disputed EEZ borders, international traffic and the concentration of critical infrastructure around chokepoints create environments where hostile actions can be easily disguised as standard operations. This has reinforced the need to develop surveillance systems that incorporate anomaly detection and attribution systems capable of distinguishing hostile behaviour from background noise.

Current techniques in maritime hybrid operations aim at creating disruptions below the threshold for armed conflict, often exploiting cost to benefit ratios that make responding to threats unprofitable. In order to distort the operational picture, adversaries use techniques such as AIS spoofing, electronic warfare, cyber warfare or, on a larger scale, disinformation campaigns. By doing so, malicious actors obscure their actions and significantly raise the risk of responding decisively, shielding themselves behind the risk of false positives resulting in conflict escalation.

As a result, maritime intelligence systems are evolving from reactive monitoring into predictive algorithms capable of aggregating multiple fragmented data sources to reduce ambiguity and support strategic decision making.

2.4.2. Maritime Chokepoints.

Maritime chokepoints represent areas of high risk and have been the proving ground for the development of current hybrid warfare techniques as seen by examining the current state of hybrid operations. Some of the most notable examples include the Strait of Gibraltar, the Strait of Hormuz, the Bosphorus of the Strait of Taiwan, and act as geographical bottlenecks for global trade, energy supply chains and military convoys. Due to the concentration of maritime traffic within narrow corridors, disruptions in these regions can generate disproportionate consequences, acting as a multiplier for offensive cost benefit ratio.

Additionally, the vessel density creates a complicated network of interaction between them, making them more difficult to monitor due to the volume of overlapping vessel activity. This complex environment aids the development of strategies since zone operations target environments that favour ambiguity, helping malicious vessels to blend into regular traffic, resulting in chokepoints being the perfect development sandbox for hybrid tactics.

The growing dependence of modern economies on global interconnection has further increased the strategic importance of chokepoints, as they are often the optimal path for these interconnections. As an example, the Suez Canal is transited by around 15% of global trade, with the Red Sea hosting 17 submarine cables handling data connecting Europe, Asia and Africa. These cables handle approximately 90% of the data traffic between Europe and Asia, almost guaranteeing a collapse of the global digital economy if they were to be severed. As a consequence, it is crucial for maritime intelligence systems to include infrastructure defence systems, with a special focus on chokepoints

as the development ground for tactics and the most threatening theatres of operations.

2.4.3. Attribution, Deterrence and Escalation Dynamics.

Since hybrid operations rely on shaping escalation dynamics and manipulating thresholds to increase the scope of operations that can be executed, it is important to maintain a grounded perspective and avoid allowing scope creep. Hostile actors attempt to maintain plausible deniability by mechanisms such as delaying escalation, acting through proxies, complicating political response or hiding behind civilian infrastructure. All of these techniques attempt to separate the response from the action as much as possible, facing a softer response that sets a precedent. Furthermore, shrouding operations in ambiguity also allows adversaries to probe responses without risking significant political capital. In practice, many responses are subdued not because the challenge is not detected but rather because generating sufficient evidence to support confident attribution and a weighted response is often difficult.

One of the main ways of facilitating confident responses is by gathering decisive evidence, assisted through the use of techniques such as integrating data fusion into intelligence solutions. By aggregating different sources, systems become capable of correlating behaviour across a variety of domains, substantiating accusations. The implementation of these frameworks represents a paradigm shift, reversing the attribution dynamic. Previously, exposing grey zone operations required an overwhelming amount of confidence in a shallow data pool. Because data streams were uncorrelated, any missing link or signal error allowed adversaries to weaponize the ambiguity, dismissing evidence as inconclusive. However, instead of requiring certainty from one sensor, modern architectures rely on multidimensional sensing systems that compare collected data against standard baselines. If even a minor signal fails to align with the expected vessel profile surveilled units can be flagged as suspicious. The breadth of the sensing systems serves as the main advantage in detecting adversarial activity,

as perfectly sanitising and keeping OPSEC across every data system simultaneously means that hybrid operations are significantly harder to hide and easier to uncover, gathering definitive proof in the process.

However, it is important to be mindful of the effect of automated systems in the political and strategic environment, especially because of the complexity of escalation dynamics and the grave consequences uncontrolled escalation can have. If not presented correctly, behavioural classification systems my increase political pressure for immediate responses in uncertain environments, meaning that maintaining human oversight is crucial to the implementation of such systems, also ensuring that they function as an aide to human decision making.

2.4.4. Technological Sovereignty and Intelligence Dependence.

Increased global connectivity and the digitalisation of intelligence systems have also raised concerns about system ownership and technological sovereignty. Modern intelligence systems depend on factors that are concentrated within a limited number of states or corporations such as cloud infrastructures, semiconductor supply chains and satellite constellations, which could become vulnerabilities during geopolitical crisis.

For maritime intelligence systems, dependence on external infrastructure introduces risks such as data access restrictions, service denial or foreign influence over compute resources, which can weaken strategic decision making resources as well as being an OPSEC risk. Many cases of operational security failures and strategic espionage are a result of improper procurement of both hardware and software, which creates systemic vulnerabilities at a deeper level than any hacking attack. This is especially relevant with countries applying military-civilian fusion laws that require companies and citizens to support and assist intelligence agencies in their tasks. An example of improper procurement creating vulnerabilities can be taken from the implementation of the European Union 5G network. Through the late 2010s and early 2020s, telecommunications companies tasked with building the European 5G

network relied heavily on subsidized Chinese equipment from Huawei and ZTE, both active participants and heavily involved in the Military-Civilian Fusion strategy. Recognition of these vulnerabilities led the European Commission to issue a formal recommendation on the 4th of May 2026 to replace any telecommunications infrastructure provided by Huawei and ZTE. Although it is not legally binding, Henna Virkkunen, the EU's Vice-President for Technological Sovereignty presented a cybersecurity package on the 20th of January 2026 as a step to convert the softer recommendation into a formal ban on sovereignty-threatening imported technology. One policy aspect that is especially relevant to the case at study is the expansion of the recommendation to "connectivity infrastructure" rather than 5G networks and radio access, the main focus point of previous policies. This comes with a broader scope of European measures to include submarine cables and satellite networks.

This trend nor only applies to software, cloud infrastructure or hardware for the systems themselves, expanding as national security becoming increasingly intertwined with industrial policy. As intelligence architectures become more dependent on access to advanced computational resources, secure energy, raw materials and manufacturing capabilities has become a strategic priority, allowing for the development of tools and resources to guarantee national security through a wider plan for strategic autonomy.

2.4.5. Intelligence Saturation and Decision-Making Risks.

A paradox developing in the modern intelligence environment is the decorrelation between collection capability and decision-making quality. While increased collection capability was the main development focus before the digital era, the overwhelming access to data from a variety of sources means that an increase in situational awareness capabilities introduces the risk of intelligence saturation, where analysts are exposed to data volumes that exceed their processing capacity. Another danger of increased data collection

comes from the natural decision making process followed by analysts. As explored by Richard Heur in his book, Psychology of Intelligence Analysis, an increase in the items of information factored into a decision not necessarily increased the accuracy of the interpretation but significantly increased the analyst's confidence in the decision, leading to a false sense of security.

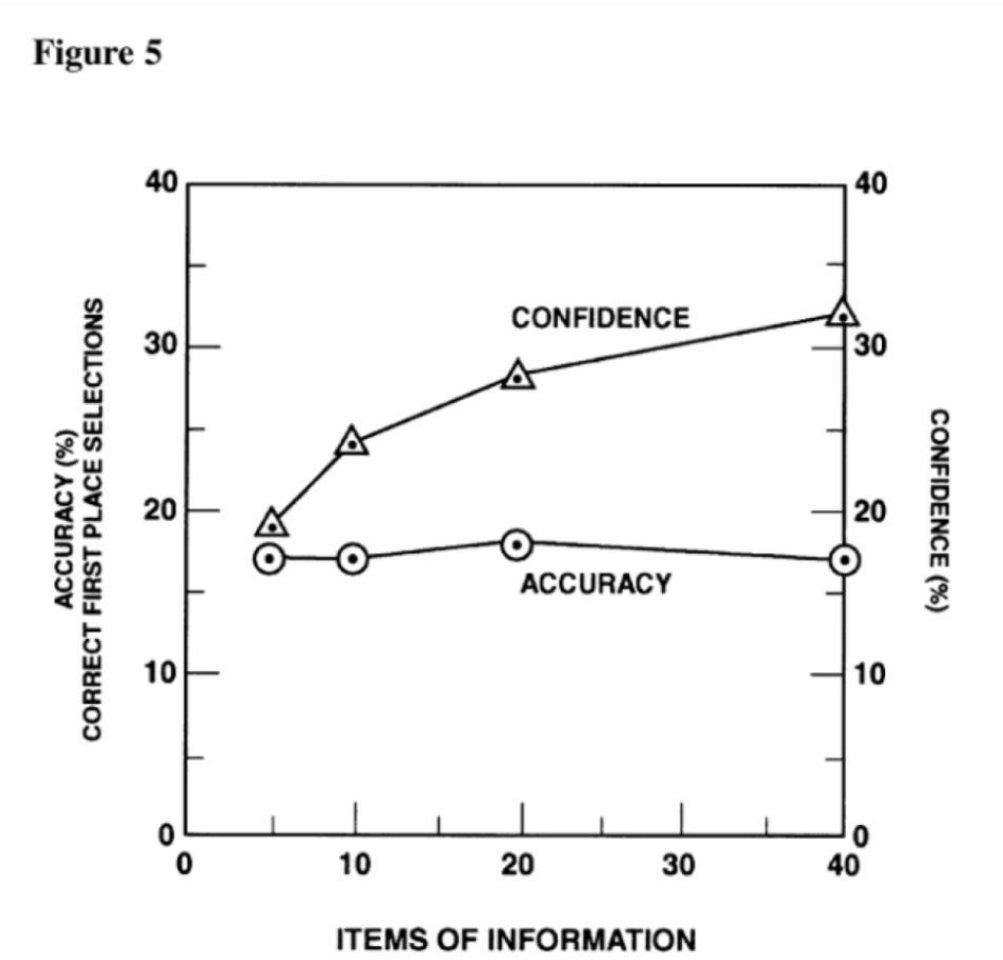

Fig. 17. Accuracy and Confidence measures for a decision as a function of Items of Information received. Richard Heur, Psychology of intelligence analysis.

This mismatch between confidence and accuracy is reflected upstream through the data pipeline and could lead to strategic blunders or arguments for aggressive escalation. Furthermore, it presents an opportunity for grey zone actors to exploit this cognitive vulnerability, feeding high volumes of disinformation across multiple channels. By utilising data filtering and

processing systems, the number of items of information reaching the analyst can be distilled into highly verifiable, correlated analysis packets, avoiding this dilemma.

However, excessive dependence on automated systems also has its dangers, which is why "human in the loop" systems are crucial. One of the risks of adopting AI decision making platforms comes from the acceleration of the decision loop, which although a seemingly positive trait can have a negative strategic impact. One of the main decision loop models is John Boyd's OODA loop model, derived from his experiences as a US Air Force pilot during the Korean War. The decision loop model is comprised of four sequential stages that form a continuous loop. The first or "Observe" stage refers to the data gathering process, the second "Orient" stage is applied as the data flows upstream and is analysed. In the second half of the loop a course of action is decided upon based on the orientation (the "Decide" phase), with the last phase referring to the execution of the decision (or "Act" stage). It is important to note that only the first two stages of the loop can be accelerated in accordance with the proposed model constraints, requiring human involvement to decide on strategic action. Boyd's theory states that by accelerating the OODA loop, it is possible to "operate within" the adversary's loop, leaving them unable to react. However, this increased pace to operate at machine speed can also lead to an unstable command environment, where human decision makers are forced to abandon slow methodical analysis and adopt reactive thinking, leading to a fast, automated, escalation. Furthermore, high volumes of low confidence alerts can also fatigue analysts, reducing the effectiveness of automated systems and slowing the decision phase of the cycle.

As a result, modern intelligence systems are shifting from data gathering exclusively towards information filtering and interpretability. Explainable AI, a concept developed by DARPA referring to AI algorithms that produce easily interpretable, transparent justification for their outputs, is becoming increasingly important for this reason. This becomes crucial when analysing

hybrid operations, where opaque model reasoning acts as an ambiguity multiplier in an already confusing environment.

2.4.6. Operational and Strategic Challenges.

One of the principal difficulties when countering hybrid maritime activity is the operational complexity of the environments in which they are conducted, kept so intentionally for maximum effect. Hybrid campaigns deliberately exploit regions with overlapping jurisdictions, dense civilian traffic and ambiguous jurisdiction and governance structures, creating situations where identifying hostile intent from anomalous behaviour becomes challenging.

This complexity is especially visible in strategic chokepoints such as the Strait of Gibraltar, where the coexistence of multinational interests and private shipping concerns intersect at one of the densest maritime traffic corridors in the world. Within the Strait, state actors such as Spain, the UK (through its presence in Gibraltar), Morocco, Algeria (through its conflict with Morocco and ownership over gas pipelines), the United States and Russia or China all maintain distinct, sometimes competing strategic interests. While Spain, the UK and Morocco are all directly involved through territorial access, other external actors have different reasons of interest in the theatre, ranging from infrastructure protection to trade security.

For example, Algeria is heavily invested through its conflict with Morocco and the transit of the Medgaz pipeline, with the United States being involved through its naval base in Rota and the ability to conduct maritime surveillance operations as well as avoiding creating a power vacuum for adversaries to fill. Two actors willing and actively attempting to take advantage of the situation are Russia and China, the first through its indirect interests in the Mediterranean (Including supporting proxies and PMC operations in North Africa through its hybrid fleet and the Africa Corps), as well as through the expansion of their ghost fleet activities and the erosion of western influence in

North Africa. China shares some of Russia's objectives with the severing of western influence in North Africa for the development of the Belt and Road Initiative, including the exploitation of the Tangier Med Port in Morocco and the El Hamdania Port in Algeria, as well as strategic ports on the European coast such as the influence in the port of Valencia.

In addition to the state dynamics there are organisational interests at stake, with NATO, the European Union and the Arab League all forwarding their interests in topics such as border control and security (with EU involvement through FRONTEX and the Spanish Civil Guard), energy and trade flows between the European and African continents or, in the case of NATO, to facilitate the monitorisation of one of the most complex chokepoints to ensure trade stability. The Strait of Gibraltar also proves to be especially important to NATO as a connection between EUCOM and the exclusively American AFRICOM. Operating in parallel with this environment, non-state actors such as trafficking networks, smuggling organisations or local proxies aiming for regional destabilisation also operate within the same theatre, often using commercial traffic as concealment for illicit grey-zone operations.

The fragmented governance structure of chokepoints like the Strait of Gibraltar are another reason for the increased operational uncertainty. Border disputes or EEZ delimitation issues and differences in jurisdiction, legal interpretation or even data access and surveillance coverage can create uneven responses which grey zone actors could actively exploit. Incidents occurring across multiple territorial zones may be interpreted differently depending on the parties involved, making deterrence more difficult. Furthermore, given the sensitivity of border politics across neighbouring countries, taking action without sufficiently robust evidence as a backing may risk diplomatic escalation, meaning that intelligence systems must be aware of the impact of false positives.

Another challenge is the evolution of adversarial behaviour as grey zone tactics develop. If actors become aware that certain behavioural patterns or vessel profile are more susceptible to triggering surveillance systems, they

may easily modify their operating procedures or inject misleading signals into the operational environment, creating a continuous feedback loop. As a result, effective intelligence systems will not only be required to detect anomalies, but to anticipate behavioural adaptation over time through models that evolve as threat profiles do.

2.4.7. Future Trends in Hybrid Operations.

The rapid development in hybrid operations is expected to continue in coming years as they represent a change in the structure of international conflict. Because traditional near-peer conflict carries inadmissible risks for most modern societies, adversarial states have optimised their strategies to persistently operate below the threshold for legally defined warfare. This “Soft War”, as coined by theorists such as Michael Wilson, is designed to coerce or degrade an adversary through unarmed tactics, including cyber warfare, economic sanctions and blockades, lawfare or Information warfare (IWAR), a match for modern hybrid operations. As they operate at the forefront of technology, it can be assumed that they will follow trends in technological development and become increasingly decentralised, automated and across multiple domains, with future conflict likely relying less on direct kinetic (hard) confrontation and more on exerting pressure through informational, economic and technological means.

A particularly relevant emerging theory is proposed by Space Force Guardian and National Defense Fellow Jason P. Lowery in his book, SoftWar, where he argues that modern competition is shifting towards non-kinetic power projection within a digital ecosystem. In this environment, deterrence is framed as power projection, defined as the ability of a state or actor to impose physical, economic, or computational costs on an adversary, Lowery also characterises systems capable of resisting disruption as having “natural proof of power”, imposing cost by causing the wasting of resources and energy (measured in Watts). While this theory was originally developed around the cyber domain, its principles are easily transferred to hybrid maritime operations and their environment, shaping how awareness and command and control structures are developed. It is also important to note that, with Lowery’s appointment to Special Assistant to the Commander

of INDOPACOM (Admiral Samuel Paparo) in June 2025, this framework has moved beyond the theoretical level, being implanted in the theatre as a way to use digital protocols such as bitcoin as a method of power projection.

One of the ways in which these protocols could be implemented into defence against grey zone activity is through the use of security protocols to stabilise chaotic, constantly evolving environments. The SoftWar framework defines strategic competition as a "struggle against entropy", in which states attempt to maintain the structural stability, organisation and integrity of their own systems while degrading adversarial ones. Maritime grey zone operations closely mirror this behaviour, achieving its objectives through GNSS/GPS spoofing, AIS sabotage or destruction of critical infrastructure among others. All of these attacks are designed to insert entropy into the system, informational entropy in the case of corrupting data pipelines. By integrating security protocols into sensing networks or cloud nodes for instance, abstract digital data can be linked to the "real", physical world. For example, requiring proof of work in the form of a token or access key to access sensing platforms would physically raise the cost of action (CA), lowering the Benefit Cost Ratio of Action (BCRA) that characterises grey zone activity. This in turn would also reduce system signal noise, with any suspected sabotage considered to be intentional as a conscious effort would be required to degrade system performance.

It is also becoming increasingly evident that the challenge is no longer tied to any single physical domain. Rather than being defined by the environment in which they occur, modern hybrid campaigns are characterised by the exploitation of complexity inherent to their operating theatre, as well as difficulty in attributing hostile intent or information ambiguity. As a consequence, future intelligence architectures must be increasingly capable of operating across multiple domains while applying common analytical principles to identify and react to suspicious behaviour. This suggests that future hybrid operations will not remain confined to traditional maritime, land or cyber domains but will instead gain relevance wherever critical infrastructure and economic interests become sufficiently valuable and interconnected. While the latest developments in grey zone activity have been linked to the maritime domain, defence trends suggest that space will

be the next primary theatre of operations for aggressive hybrid campaigns. In order to understand the translation of operations between both environments, it is first essential to highlight the similarities between space and maritime domains. In first place, both are shared global environments or "international commons", and must remain open to global commerce without restrictions as defined by both UNCLOS and the Outer Space Treaty, which state that outer space and the High Seas are not subject to national appropriation "by claim of sovereignty, use, occupation or any other means". In this regard, both domains also share a similar governance structure defined by international treaties, which share additional similarities like the Flag-State jurisdiction, which establishes that the launching state retains absolute control over a space object registered in its name, much like a vessel at sea is considered an extension of the sovereign territory under whose flag it is registered. Another shared characteristic is the unforgiving nature of both environments which leads to the development of rescue agreements such as SOLAS for maritime rescues and the Rescue Agreement of 1968 for space, which states that astronauts must be regarded as "envoys of mankind" and therefore must be offered all possible assistance in the event of an accident.

Beyond these similarities, the maritime and space domains also show physical parallelisms that make them unique to hybrid operations. For instance, analysts are heavily reliant on mediated reality technology to observe the theatre of operations, relying on sensor networks, digital communications and online telemetry to create an operational picture. An example of this would be the use of AIS transponder data or Space Domain Awareness radar to enable accurate command and control. This is due to the physical isolation of both environments, which creates a shroud of plausible deniability for hostile actors. Another interesting resemblance can be drawn between maritime chokepoints and orbital slots, becoming increasingly crowded as satellites are progressively launched. Applied to space operations, highly demanded orbits such as sun synchronous bands and polar corridors for low Earth orbits form chokepoints that contain most vital satellites, including for communications and intelligence.

Grey zone activity can also be inferred by examining the development of current satellites and space operations, with supposedly scientific experimental designed

coordinating to execute Rendezvous and Proximity Operations (RPOs) under the cover of civilian missions. One such example is incident highlighted my General Michael Guetlein, Vice-Chief of Space Operations, in which three People's Liberation Army Aerospace Force Shiyan-24C and two Shijian-6 05A/B satellites conducted "space dogfighting operations". Another concerning development is the evolution of satellite design towards more manoeuvrable designs, including satellites such as the Shijian-21, with robotic arms capable of grappling objects in space or the Jackal spaceship, designed to for future high tempo engagements. Orbital flexibility has facilitated the development of satellite-to-satellite intelligence gathering, both through SIGINT platforms and optical imaging, often basing approaches on classic dogfighting techniques such as sun-masking for concealment or aggressive deceleration manoeuvres. Other similar maritime tactics applied to current space warfare include signal jamming and spoofing, creating scenarios like the jamming of GPS navigation for thousands of airplanes in the Baltics[42] or European Commission President Ursula von der Leyen's plane in September 2025[43].

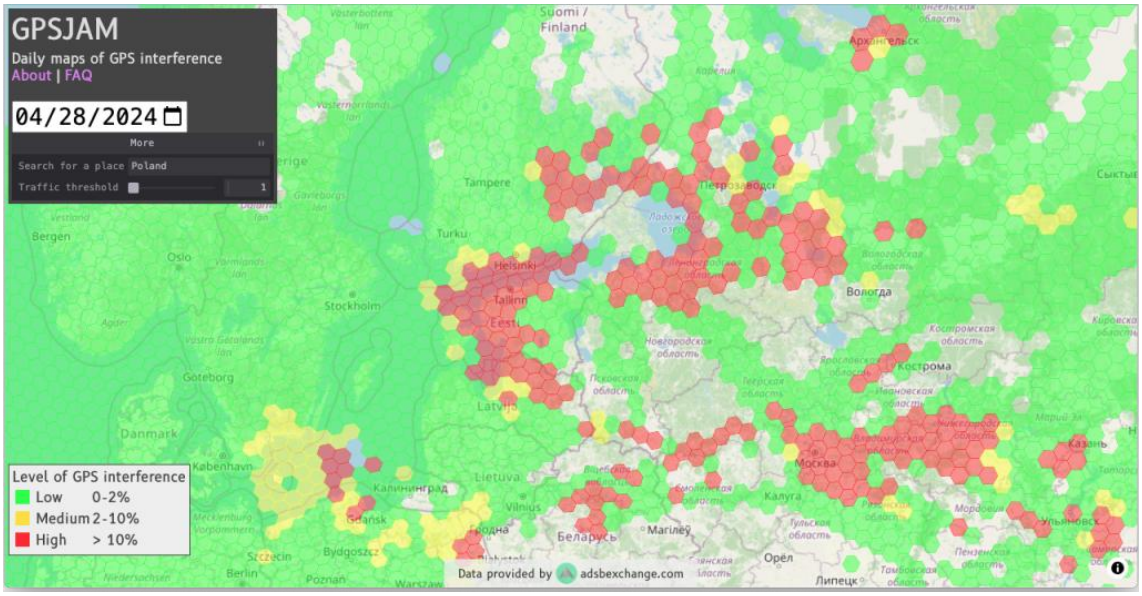

Fig. 18. Map showing GPS interference in the Baltic region, produced by John Wiseman, broadcasted on GPSJAM.org.

[42] Russian Jamming Wreaks Havoc on GPS. Is It Hybrid Warfare?
[43] GPS jamming on von der Leyen's plane the latest sign of Russia's hybrid warfare

Other current developments include the research and effort invested into the development of stealth satellites, either through the use of infrared absorbent layered materials[44] to camouflage satellites against the background sky or by minimising the Radar Cross Section (RCS) of microsatellites, such as the Olive-B.

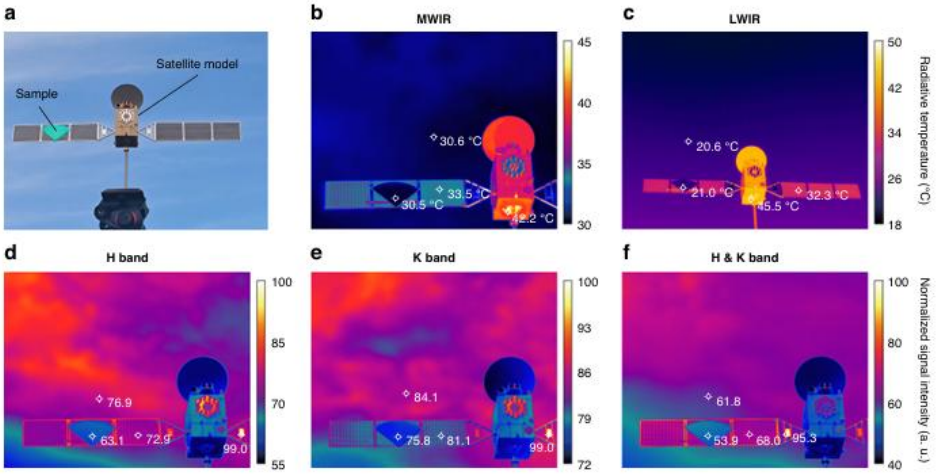

Fig. 4 Demonstration of space-to-ground infrared camouflage (the H, K, MWIR, and LWIR bands). **a** The sample is mounted onto the exterior of a satellite model, positioned outdoors in an open area. **b**, **c** The MWIR and LWIR band infrared images of the sample and the satellite model, in the form of radiative temperatures. **d**–**f** The H band, K band, and H & K band infrared images of the sample and the satellite model, in the form of normalized signal intensity

Fig. 19. Infrared camouflage applied to satellite solar arrays. Qin, B., Zhu, H., Zhu, R., Zhao, M., Qiu, M., & Li, Q. (2025). Space-to-ground infrared camouflage with radiative heat dissipation. Light: Science & Applications, 14(1), 137. https://doi.org/10.1038/s41377-025-01824-y

[44] Space-to-ground infrared camouflage with radiative heat dissipation | Light: Science & Applications

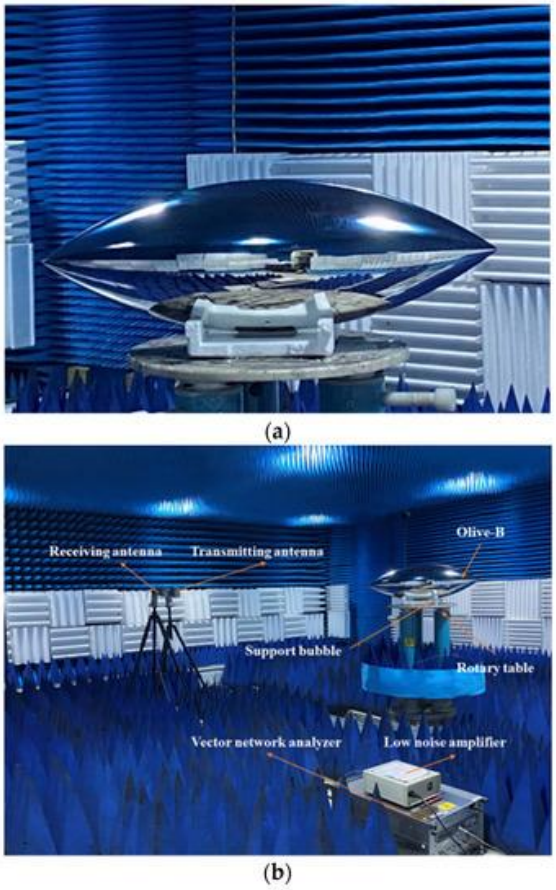

Fig. 20. RCS test of the Olive-B stealth satellite in an anechoic chamber. Sun, H., & Qin, Y. (2022). Stealthy Configuration Optimization Design and RCS Characteristics Study of Microsatellite. Aerospace, 9(12), 815. https://doi.org/10.3390/aerospace9120815

Ultimately, the convergence of maritime and space hybrid operations demonstrates that the future of conflict will depend on the ability to shape interconnected environments rather than control physical territory. As operational theatres become increasingly saturated with sensing and autonomous systems, the distinction between civilian and military infrastructure are expected to become increasingly blurred, meaning that cloud infrastructure and digital support systems can be considered crucial strategic assets and potential targets. Maintaining situational awareness in this environment is a requirement for detecting hostile activity within a dense, ambiguous environment under constant development, coming at an increasing cost as complexity increases. As a result, the effectiveness of intelligence architectures will depend not only on sensing capability but on the ability to integrate heterogeneous data streams using fast processing interpretation

systems while preserving sensible decision-making structures to reduce ambiguity faster than adversaries can exploit the gaps.

2.5. Capability Gaps.

One of the critical limitations within the current intelligence environment is the growing volume of available information. Maritime monitoring through different sources such as satellite imagery, radar systems, shipping databases or AIS transmissions can lead to an improved situational awareness if handled correctly, although it can lead to an "information overload" issue if analytical capacity doesn´t increase at the same rate as the data stream volume. The challenge is therefore not only collecting relevant data, but prioritising and interpreting it effectively to prevent losing significant data in the constant stream of routine maritime activity. Applications capable of integrating information from multiple sources into a unified analytical environment through data fusion could significantly alleviate the problem by filtering large data streams and highlighting the most urgent threats based on predictive mechanisms. By presenting analysts with prioritised behavioural indicators instead of raw datasets, similar systems could help transform large volumes of data into actionable intelligence.

The increasing mobility of hybrid threats also contributes to widening existing capability gaps. Modern maritime operations are no longer constrained to specific regions of interest, due to the globalisation of operating theatres and the ease of access between them. Recent activity involving Russian vessels operating across routes linking the Persian Gulf, the Caribbean and the North Atlantic over a couple of days highlights this issue[45]. This mobility complicates the identification of behavioural patters when the monitoring is conducted within isolated geographic areas, for this reason, comparative case studies across different maritime regions have become an important analytical tool, allowing patterns observed in one theatre to be used as training for monitoring strategies being applied to others.

In response to these emerging challenges, private maritime intelligence companies have recently started to develop platforms designed to address the problems at hand. The most complete and relevant example as of March 2026 is Starboard Maritime Intelligence, a platform that combines vessel tracking data with satellite imagery, radar inputs and other sensor streams to provide a detailed picture of maritime activity. These systems apply machine learning models to the information collected to identify anomalous behaviour or assess potential risks, and do so from a vast amount of collected information, with some platforms analysing more than one billion data points per day [46]to identify threats and highlight suspicious activity.

Despite these technological advances, many existing platforms still focus primarily on improving situational awareness instead of developing analytical frameworks capable of

---

[45] https://www.reuters.com/business/energy/us-seizing-venezuela-linked-oil-tanker-after-weeks-long-pursuit-2026-01-07/

[46] Starboard Maritime Intelligence secures NZ$23M Series A to scale global AI-powered maritime intelligence technology | Starboard Maritime Intelligence

interpreting intent or supporting attribution, leaving risk assessment on a strategic level to depend heavily on manual analysis. As hybrid operations adapt to become more sophisticated, the development of analytical tools capable of moving beyond anomaly detection towards predictive behavioural modelling remains an important research field.

Chapter 3. METHODOLOGY.

The project follows a structured design and development sequence organised around three phases, each building on the outputs of the previous one.

The first phase consists of a review of the current state of affairs in hybrid operations and the technological infrastructure supporting modern intelligence systems, as shown in chapter 2, “State of Affairs”. This section is structured around identifying capability gaps in modern maritime frameworks and analysis the technical and analytical tools that are available to address them. The capability gaps to address are then explained in detail at the end of Chapter 2.

The second phase covers the design and development of the system concept. Working from the identified gaps, the architecture is built around three core requirements established in Chapter 1, the use and integration of open-source data, a strong focus on explainability at every layer and live data ingestion for near real-time operation. The JDL fusion framework provides the structural backbone of the system, with each layer of the pipeline corresponding to a level of fusion, from raw inputs being ingested at level 0, through feature extraction at level 1 to a multi classifier scoring system at JDL2. The pipeline is then extended to the space domain as a parallel case study, demonstrating that the analytical framework generalises beyond maritime surveillance without requiring a redesign of the system architecture.

The third phase addresses the validation and evaluation of the model. In this phase, a validation framework is built to determine the robustness and accuracy of the intelligence pipeline. Developed independently of the systems built in Chapter 4 to avoid destabilising the working systems, this tool produces quantitative performance metrics for both systems by generating train/test datasets to evaluate model performance. The results are assessed against the original design objectives in Chapter 5, alongside with policy recommendations and a review of model limitations and future development opportunities.

Chapter 4. SYSTEM CONCEPT.

4.1. Design Philosophy.

The increasing complexity of maritime structures requires a different approach from intelligence analysts and security organisations. While advances in sensing technologies have significantly improved awareness, the resulting increase in data volume has introduced a new challenge, the ability to transform large quantities of fragmented information into actionable intelligence on demand. This challenge is particularly relevant for hybrid threat analysis, since traditional monitoring systems often struggle to distinguish deliberately ambiguous suspicious activity from legal background activity, frequently relying on human expertise to help with classification.

The proposed system has therefore been designed as an AI-assisted intelligence support program, designed to assist human analysts in confidently assessing threats rather than replacing them, otherwise known as human in the loop architecture. To achieve this, AI assistance is constrained to the tactical rather than strategic data levels in the JDL framework, with the AI layer enabling faster data ingestion on a larger scale, pattern identification and the presentation of evidence-based assessments, allowing for human operators to make strategic decisions based on transparency.

The developed architecture is consequently based on three principles, explainability, human oversight and awareness of uncertainty. In first place, enforcing system explainability ensures that analysts can understand the factors contributing to the generated assessments and confidence metrics, key in developing a solid base for human analysis. Human oversight guarantees that AI generated recommendations are solely advisory and subjected to review, permitting explainability at a strategic level, with decisions being backed by human assessment and reasoning. Lastly, by recognising the inherent complexity of the environment, the system assures that confidence metrics reflect the

operational environment, establishing boundaries to avoid overconfidence and false positives.

4.2.System Objectives and requirements.

The primary objective of the proposed system is to provide an environment capable of supporting the detection, assessment and investigation of anomalous vessel activity within areas of interest. Primary users benefiting from such aids would be intelligence analysts at a national and international level, including organisations like NATO MARCOM or coast guard organisations like FRONTEX. However, the growing interest in maritime awareness and security research for private businesses also opens the possibility of expanding the user base to any sector requiring maritime awareness where large quantities of heterogeneous information must be analysed and interpreted in real time.

To achieve this objective, the system must perform several functions. In first place, it must process vessel tracking information obtained from sources such as AIS. AIS data provides the foundational representation of vessel behaviour upon which analytical processes can be built. This can be achieved through transforming raw inputs into vessel trajectories, extracting characteristics like velocity profiles, heading, route deviations or transmission gaps. The system must also integrate contextual intelligence from open sources. These sources may include vessel ownership records such as the International Telecommunication Union's database[47], sanctions databases, corporate registries or historical intelligence assessments. By combining behavioural observations with contextual information, the system provides a more accurate representation of maritime activity. This can also be achieved by using machine learning techniques to establish behavioural baselines and identify anomalies that may indicate the presence of unusual or suspicious activity. However, it is important to clarify that the detection of anomalous behaviour does not imply malicious intent, serving instead as an indicator requiring further investigation.

[47] Ship Station List

Lastly, the system must provide analysts with an intuitive way to visualise analysis results, with information being presented in a manner that supports quick situational awareness. This includes the visualisation of vessel trajectories, ownership networks, threat indicators, uncertainty estimates and supporting evidence.

4.3. System Architecture.

The developed prototype applies the philosophy of explainable AI to maritime threat assessment, as well as the human in the loop design while following a layered approach inspired by the JDL data fusion framework. In this applied framework, level 0 refers to the initial data ingestion and normalisation which is followed by the refinement and feature extraction in level 1 (L1). Lastly, the outputs are interpreted using various models in level 2, being displayed in a dashboard focused on supporting strategic decision making without replacing it. The code and implementation was carried out with the assistance of Claude Code (Anthropic) and Cursor.

In order to accurately evaluate vessel behaviour, level 2 uses three machine learning classifier models, a Random Forest, an Isolation Forest and a gradient boosted XGBoost classifier, whose outputs are fused into a single probabilistic threat score. A Kalman filter based trajectory prediction model also runs in parallel to forecast geofencing breaches as explored in section 2.3.2.2. The system also runs a behavioural analysis module to determine if kinematic and contextual data into hypotheses reported in alignment with NATO intelligence reporting standard STANAG 4559. This is particularly important as it allows a common reporting system based on two principles, the first being the use of an indicator based reporting structure, with intelligence assessments being assigned a confidence level individually. This is stated in STANAG 4559 as the need for every intelligence product to identify the observables that support assessments, separating factual observations from inference. The second principle refers to the need for intelligence products to be structured around competing analytical hypotheses, each followed by their assigned confidence rating, which is done in

the model by assigning probability scores for 5 different vessel states and normalising the values on a 0 to 1 scale, reporting the confidence level for each vessel state and highlighting the margin between the top two possible states. In order to reinforce the system explainability, every confidence score highlighted in the dashboard is also designed to be accompanied by the quantified contribution of each input feature (calculated using the Shapley Additive exPlanations, or SHAP model, merging game theory and machine learning to determine the average marginal contribution of each feature across all possible combinations to come up with a final "contribution value").

4.3.1. Data Fusion and Feature Extraction. JDL 0.

Raw AIS messages are received via a livestream connection to open source data from AISStream[48] as a primary source through the API, defined by a bounding box covering the Strait of Gibraltar. Additionally, AIS data is taken from AISHub[49], VesselFinder[50], MyShipTracking[51] and BarentsWatch[52] as secondary sources as a back up for when data is sparse, making the system more robust. Vessel registry information is retrieved from the International Telecommunication Union's Maritime Mobile Service Identity Database (ITU MARS[53]), the international registry for MMSI assignments and registered ownership, with data being accessible via MMSI lookup and returning information such as ship name, call sign, vessel flag and administration code, owner name and address, IMO number, vessel type classification, gross tonnage and former ship name, with this field being especially interesting as frequently renamed vessels are highly suspicious of covert action, transforming this field into a primary input of the SHAP scoring system.

---

[48] aisstream.io

[49] Free AIS vessel tracking | AIS data exchange | JSON/XML ship positions

[50] Ship & Container Tracking - VesselFinder

[51] My Ship Tracking Free Realtime AIS Vessel Tracking Vessels Finder Map - ship tracker

[52] Front page

[53] Ship Station List

With regards to sanctioned vessels and countries, maritime sanctions lists are accessed via lookup on the OpenSanctions[54] platform, with additional risk being assigned to specific vessels depending on the flag risk associated to them. This is done through the use of three sources, the FATF black and grey lists[55], returning different levels of risk (0.98 for black flags and 0.58 for grey flags) depending on the country' status and sanctioned state as of February 2025, when the list was compiled. The system also uses the Paris MOU port state control performance list[56], which conveniently uses the same black, grey or white metric which is assigned to countries as a function of the frequency of port state control detention rates and the use of country flags in recorded hybrid operations and as flags of convenience. Finally, the complete 2024 ITU Maritime Identification Digits table is used to determine a vessel's nationality directly through the use of its MMSI without the need to perform a lookup, allowing the system to load before the ITU MARS check is completed. Port destinations are also compiled into a list that allows the system to flag high risk ports based on sanction regimes or recorded hybrid operations as an input, with examples of high risk ports including Rason in the DPRK, Bandar Abbas in Iran, or Kavkaz in Russia. Finally, environmental data is drawn from the Agencia Estatal de Meteorología (AEMET) by querying Tarifa's meteorological station (ID 5996) and the Gibraltar maritime zone (area ID 64) to draw observations feeding into the Kalman filter's noise covariance. This allows the filter to scale the process noise factor depending on the wave height, period, direction, wind speed and direction, all contributing factors to movement uncertainty.

[54] UN Security Council 1718 Designated Vessels List - OpenSanctions

[55] "Black and grey" lists

[56] White, Grey and Black List | Paris MoU

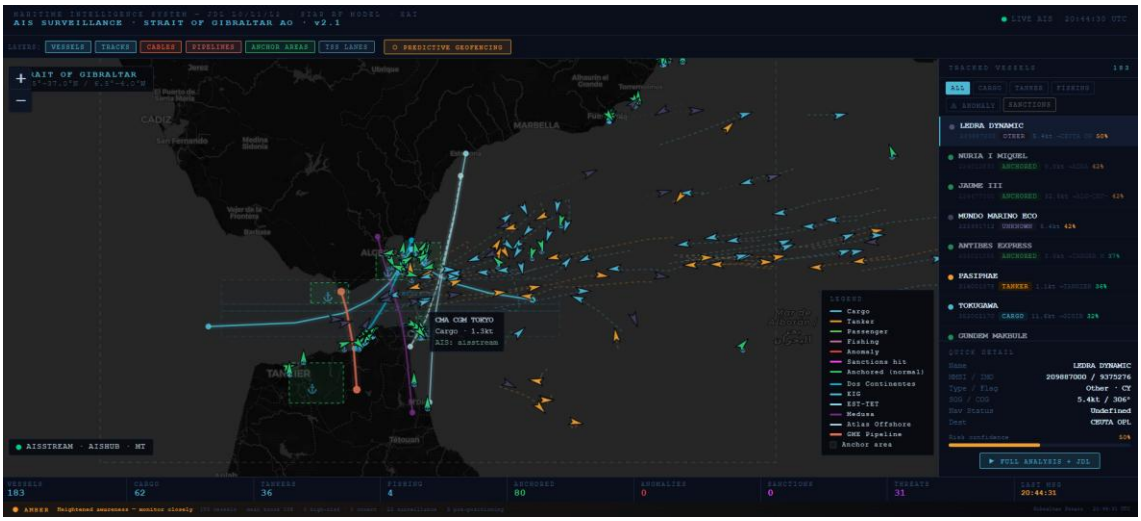

Fig. 21. Screenshot showing the active maritime anomaly detection system, showing live AIS-tracked vessels in the Strait of Gibraltar and listing tracked vessels on the right of the screen.

In order to detect outliers and data anomalies, such as implausible updates and corrupt records, a haversine distance check is performed to determine if updates are grounded in reality. This is also enabled by loaded vessels keeping 80 track points and 30 speed over ground and course over ground observations along with their timestamp, being updated as new inputs are gathered. Timestamp recording also proves to be especially interesting as transmission gaps are not interpolated but instead highlighted by the system, being considered as indicators of potential covert activity.

4.3.2. Vessel Feature Vector Composition. JDL 1.

After the data gathering phase, the main features are assembled into a 13 dimensional vector, with kinematic features being extracted from the AIS track and contextual features being taken from a variety of the sources mentioned above. The features compiled in the vector include:

- Time elapsed since last valid AIS position report.
- Number of significant heading changes. (Exceeding a change in COG of over 25º)

- Speed variance, represented by the standard deviation of the SOG history track ($\sigma^2 = (1/n)\sum(SOG - \mu)^2$.
- Flag risk score.
- Exposure to sanctions.
- High risk port calls.
- Ownership changes taken from the ITU MARS register.
- Abnormal routing, a result of heading deviations relative to the traffic corridors along the Strait.
- Hours of activity in the dark, measured by the time spend in proximity to infrastructure without AIS signal.
- Cargo risk.
- Cargo volume percentage.
- Crew criminal records.
- Port state control detentions.

These indicators have been taken in part from the SIAP dataset as explained by Karamolegkos et al., 2025 in the paper "SIAP: Synthetic dataset for maritime vessel risk profiling and illegal activity prediction"[57]. In order to avoid penalising vessels incorrectly due to lack of data, unavailable records will be assigned neutral values which are zeroed when scoring.

[57] Karamolegkos, S., Dourvas, N., Episkopos, N., Pikounis, K., Michail, E., Ioannidis, K., & Vrochidis, S. (2025). SIAP: Synthetic dataset for maritime vessel risk profiling and illegal activity prediction. *Data in Brief*, *63*, 112101. https://doi.org/10.1016/j.dib.2025.112101

Fig. 22. Figure showing the Panamanian-flagged tanker "DOLPHIN 19" selected in the maritime anomaly detection prototype.

Fig. 23. Vessel identification kinematic data for DOLPHIN 19.

Fig. 24. Pre-processing and object refinement levels for vessel "DOLPHIN 19".

4.3.3. Anomaly Scoring. JDL 2.

The main behavioural classifier is a Random Forest Model trained on 40000 synthetic vessel records to match the SIAP dataset which is split at an 80/20 ratio between legitimate and suspicious vessels. The model operates by creating 120 decision trees (with a maximum depth of 8, shallower than normal to avoid overfitting, one of the main risks of using RF algorithms) which use different subsets of the training data and features. The final classification is obtained by combining the output of all trees. An RF model was chosen to be the primary analysis algorithm due to its ability to handle mixed format data, both in numerical and categorical form, as well as its capacity to process datasets with incomplete information and relatively simple format, helping with explainability.

Fig. 25. Risk and confidence score for vessel "DOLPHIN 19".

Anomaly detection on the other hand falls to an Isolation Forest algorithm, designed to identify vessels with significantly different behaviour to surrounding maritime traffic. One of the advantages of the IF model is that, unlike other classification algorithms, it does not require predefined threat categories to be assigned to the system, being trained on exclusively legitimate vessel distribution and focusing on identifying statistical outliers. This is especially relevant for hybrid operations as new grey zone tactics may not be included in suspicious training data sets. The anomaly score is calculated from the average number of decisions (or

splits/partitions) required to isolate a vessel within the Isolation Forest. Since anomalous observations are outliers, they require fewer splits to isolate and have shorter paths along the trees. The average path distance can then be transformed into a probability of a vessel's behaviour being anomalous, with the anomaly score being inversely proportional to the expected path length.

The final model applied by the detection system is the Extreme Gradient Boosting XGBoost Classifier. This model is also composed of trees with the difference being that trees are assembled sequentially rather than independently as RF for example. In this way, each subsequent tree attempts to correct the classification errors made by previous trees, correcting the residuals generating bias in the RF model making it useful as a corrective mechanism to improve accuracy when applied at the end of the data pipeline. The XGBoost model mainly takes contextual data from sanctions records or historical operational behaviour among others to uncover non-linear relationships that the previous models may have missed, another of the main advantages of the classifier. Another desirable characteristic that justifies the use of an XGBoost classifier is its effectiveness when applied to structured or tabular data since the inputs are sufficiently processed and structured at that point in the pipeline.

Fig. 26. Scores returned by Isolation Forest and XGBoost models.

The resulting confidence scores from all three models are combined by a simple weighted function, comprising the fourth layer of the JDL 2 stage.

$$Fused\ threat\ score = 0.5RF + 0.3XGB + 0.2IF$$

The weights are assigned as a function of the model nature, for example, the Isolation Forest threat score receives the lowest weight due to its unsupervised nature, making it more susceptible to false positives. On the other hand, the Random Forest prediction score receives the highest weight as the primary supervised classifier, with the XGBoost output receiving an appropriate score to have the desired corrective effect. To maintain sufficiently strict assessments in very clear high risk cases, any vessel returning positive matches for UN or OFAC sanctions lists is automatically assigned a minimum fused score of 0.8, ensuring that vessels subject to sanction regimes cannot be masked by the intentional manufacturing of "positive" indicators through skilled tradecraft. Another important characteristic is the flagging of model disagreements, highlighting any differences greater than 0.25 in RF and XGBoost

predicted risk scores, highlighting a situation of increased uncertainty and possibly incomplete information.

Fig. 27. Data fusion level.

Fig. 28. "DOLPHIN 19" behavioural analysis result overlayed against the training distribution.

4.3.4. **Confidence Quantification. JDL 2 Layer 5.**

A key feature of the explainable AI architecture is the separation of risk and confidence scores. This is crucial for interpretability as, for example, a high threat score derived from poor quality data should be treated differently from the same score produced by highly verifiable and trustworthy sources. Therefore, the system assigns a confidence score C [0,1] for every threat assessment comprised of five independent variables.

$$C = 0.3\,C_{TRACK} + 0.25\,C_{FRESHNESS} + 0.2C_{AGREEMENT} + 0.15\,C_{FEATURES} + 0.1\,C_{STATIC}$$

The track variable is arguably the most important variable for data interpretation which is why it is assigned the highest weight in the confidence formula, as without sufficient kinematic datapoints machine learning models are unable to interpret movement and intentions. Therefore, vessels with minimal data are assigned a baseline 0.3 score, increasing to 1 for vessels with more than 30 data points. Data freshness is also treated in a similar manner, with the freshness coefficient being one for messages being under one minute old and decaying to 0 for messages over 15 minutes old. As explained in section 3.3.3, the model agreement coefficient is derived directly from the RF/XGBoost threat probability score margin, with feature completeness accounting for contextual records being absent by assigning a higher score to vessels with complete records. However, as vessels are gathered through AIS tracks, kinematic data is always going to be available, leading to a minimum score of 0.4. Finally, the static field reflects the completeness of vessel static fields, particularly draught and destination, two of the most critical indicators of hybrid operations as seen in the SIAP dataset paper.

Fig. 29. Confidence parameters explained for "DOLPHIN 19".

4.3.5. SHAP Explainability.

To ensure that the system remains transparent as required by the explainable AI philosophy the threat assessments must be able to be backed by traceable evidence. To achieve this, the system incorporates an explainability layer based on Shapely Additive exPlanations, or SHAP.

SHAP is used to determine the contribution of each feature to the final threat assessment produced by the Random Forest and XGBoost classifiers. Rather than simply displaying a risk score, the system identifies which factors increased or decreased the estimated threat level and quantifies their influence. This is achieved using the TreeExplainer algorithm presented in the paper "Explainable AI for Trees: From Local Explanations to Global Understanding" by Lundberg et al. [58] The Shapley value represents the average marginal contribution of each feature across all possible combination of features (named coalition). The impact is calculated by taking coalitions that lack the feature being examined (i) and evaluating the shift in predicted value when adding feature i to the coalition. Finally, the weighted average of the marginal contribution across all coalition combinations is calculated to ensure that the feature is not assessed on an individual basis. The system computes the SHAP values for both the RF and XGBoost classifiers and produces an attribution ranking by averaging the values across both models, scaled to account for the differences between models. Each feature is then labelled as increasing or decreasing risk based on the sign of its RF SHAP value and is presented in the attribution table with its value, neutral baseline and deviation.

[58] Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2019). *Explainable AI for Trees: From Local Explanations to Global Understanding* (arXiv:1905.04610). arXiv. https://doi.org/10.48550/arXiv.1905.04610

Fig. 30. SHAP values and contribution to risk scoring against the model baseline.

4.3.6. Behavioural Analysis and Intent Classification.

The behavioural analysis model is designed to capture anomaly scoring across the analysis timeline, critical for hybrid operations since individual events may appear to be perfectly legitimate at a certain point in time but form a pattern of intent over time. To achieve this, the model captures 6 behavioural indicators over a rolling 24 hour window, which are assigned weights when assembled into a composite behavioural score. The examined indicators are the following:

- Loitering indicator. Combines SOG (a low Speed Over Ground is assigned a higher scoring) and the haversine distance in nautical miles

over the last vessel positions with scores below 5 NM being treated as a tight cluster. Weight = 0.25.

- AIS gap indicator, not only recording current transmission gaps but its temporal repetition to detect behavioural patterns. For example, a vessel with three prolonged AIS gaps in the last 24 hours scores higher than one with the same current gap but no repetition history. Weight = 0.2.
- Route deviations are scored using the relative heading along the Strait of Gibraltar's expected TSS corridors (eastbound at 070º and westbound at 250º), flagging unexpected or non-compliant behaviour. Weight = 0.15.
- Rendezvous approaches are detected through examination of the telling sequence of deceleration, and acceleration characteristic of ship-to-ship transfers. With the first and last speeds necessarily being higher than 4 knots and the middle slow approach being slower than 2 knots to avoid flagging subtle manoeuvres. Weight = 0.15
- An Infrastructure proximity and orbiting score is also computed by assigning a penalty to vessels showing sustained proximity to critical infrastructure during a 6 hour window. Weight = 0.15.
- Speed shift profiles are also examined to assign behavioural penalties to vessels manoeuvring in an aggressive manner, with the maximum score being awarded to vessels with a delta exceeding 8 knots between consecutive readings. Weight = 0.1.

When comparing the current weighted behavioural scores with the mean of the 10 previous scores, an escalation trend can be inferred, assigning "ESCALATING", "STABLE" or "DE-ESCALATING" tags to vessels, providing analysts with updating indicators, representative of a dynamic environment.

The behavioural profile obtained for a specific vessel is then passed down the data pipeline to an intent classifier that maps the six dimensional indicator to one of five hypothesis, innocent passage, loitering, surveillance, covert transfer or pre-positioning. These hypotheses were selected because they represent behaviours commonly observed during

hybrid maritime operations, and while they do not provide definitive evidence of hostile activity, they allow analysts to move beyond simple anomaly detection at a glance to begin assessing the operational significance of observed behaviour. It is important to note that the classification is rule based and does not apply learning models to ensure explainability while also tailoring the decision weights for every specific hypothesis. For example, the covert transfer hypothesis places a higher emphasis on dark activity and rendezvous behaviour while the surveillance hypothesis assigns a higher weight to close proximity to critical infrastructure such as repeated low speed operations around energy infrastructure.

The resulting hypothesis scores are normalised into a probability distribution, allowing analysts to evaluate the likelihood of each possible explanation. Rather than presenting a single deterministic answer, the system presents uncertainty by displaying the probability with each hypothesis, reinforcing the explainable AI concept and reflecting the nature of intelligence interpretation, where multiple competing explanations often remain plausible at the same time. To support analyst decision making further, the system also calculates a confidence metric based on the probability difference between the two most likely hypothesis. Large margins indicate that the available evidence strongly support one interpretation while smaller margins indicate a higher level of ambiguity and the need for caution or more information.

Fig. 31. Behavioural assesment and breakdown.

4.3.7. Ownership and Relationship Network Visuatlisation.

Vessels involved in sanctions evasion, illegal transfers or coordinated maritime activity frequently involve a structured ownership network and are rarely conducted in isolation. These ownership structures are often concealed behind shared management companies, which is why the system incorporates a network visualisation function that presents these relationships as an interactive graph.

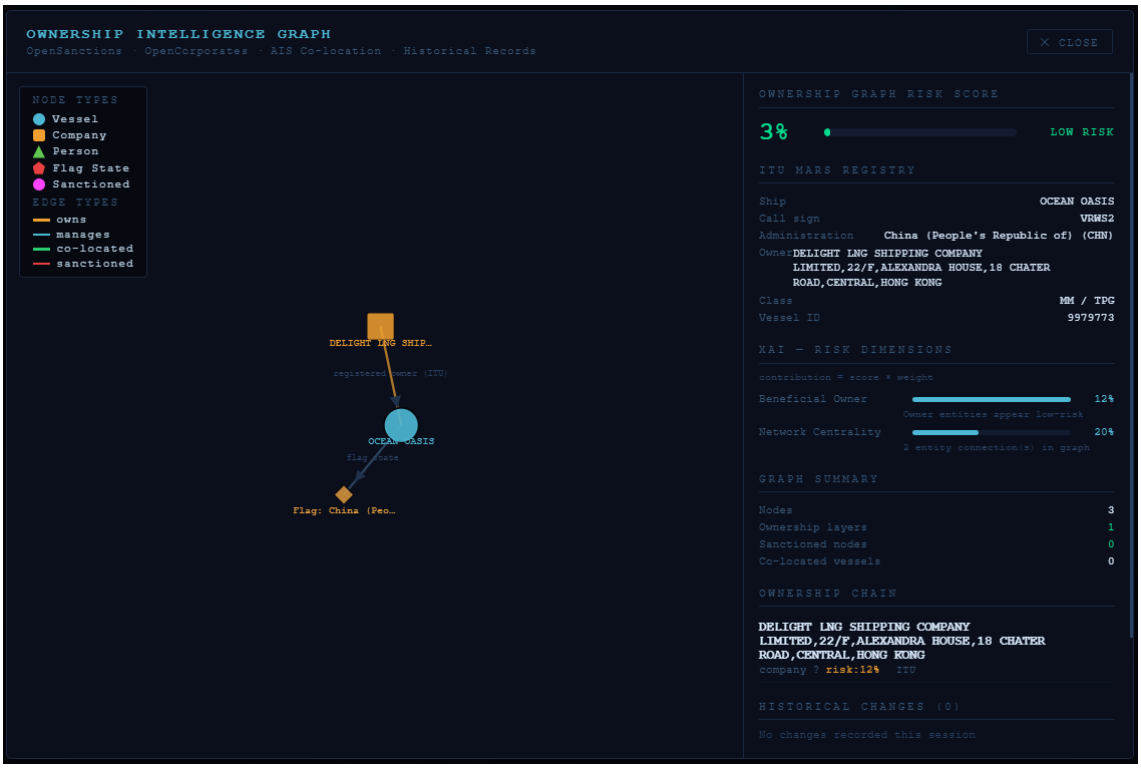

Fig. 32. Ownership intelligence graph for vessel "Ocean Oasis".

As seen in figure 32, each vessel is represented as a node within a graph, with edges created where a relationship exists between two entities. Depending on the available open-source information available, these relationships may include shared ownership, common management companies, repeated AIS co-locations or flag changes. Instead of acting like a classifier, the network serves as an analytical aid, allowing the operator to identify clusters of related vessels that may appear to be low risk but collectively exhibit patterns consistent with hybrid maritime activity.

The network is designed to be easily interpretable. Node colours represent different vessel attributes, while edges represent the relationship between two connected entities. Relationships may range from strong structural links, such as shared ownership or common management to weaker associations such as co-location events. Figure 33 illustrates this, showing the Liberia-flagged vessel “MSC CLAIRE” and the Spanish-flagged vessel “FLY BLUE TRES”. Although the two vessels share no ownership

or management structure, they are connected through a recorded co-location event within one nautical mile, providing contextual information that may support further investigation when considered alongside the AI-generated confidence score.

Fig. 33. Network ownership graph for vessel "MSC CLAIRE".

This capability aligns closely with the objectives of the JDL framework, where network visualisation assists levels 2 and 3 by exposing contextual relationships that may indicate coordinated activity. It also contributes to the explainability of the proposed system by allowing analysts to understand why a vessel may warrant further attention, complementing the confidence score with transparent evidence. The resulting view enables analysts to investigate the wider operational context behind threat assessments before initiating further intelligence analysis or enforcement actions.

4.3.8. Predictive Geofencing Through Applied Kalman Filters.

A predictive geofencing feature was also introduced into the model in order to move the system's capabilities forward from anomaly detection to threat anticipation. The objective of the developed function is the estimation of the probability that a vessel will enter a user defined zone in a user defined time zone.

The estimator is based on the model developed in section 2.3.2.2, operating on a local coordinate system centred around the vessel's last known position. The 4 dimensional state vector comprised of x and y position along the reference plane as well as the vessel's velocity vector along x and y is obtained from the AIS kinematic data, from which the transition model and process noise covariance are derived. It is important to note that the manoeuvring acceleration standard deviation, taken as a constant in section 2.3.2.2, is assigned by vessel type, being 0.01 m/s^2 for tankers, 0.018 m/s^2 for cargo vessels and 0.06 m/s^2 for fishing vessels, on the premise that the physical constraints for different vessel types vary greatly, therefore affecting the positional uncertainty over time. The process noise is also scaled by taking into account the meteorological conditions obtained from AEMET, the Spanish State Meteorological Agency and relating Beaufort scale wind observations to wave height by applying equation1.30 in the Guide to Wave Analysis and Forecasting by the World Meteorological Organisation[59], derived in turn from the Pierson-Moskowitz wave model, proving that the significant wave height (H_{m0}) for a fully grown sea is:

$$H_{m0} = 0.0246u^2$$

With u being the wind speed in m/s.

[59] World Meteorological Organization (Ed.). (1998). *Guide to wave analysis and forecasting* (2nd ed). Secretariat of the World Meteorological Organization.

AIS gaps also contribute an additional noise term proportional to the elapsed time since the last transmission, with the manoeuvring acceleration standard deviation σ_a increasing by 0.01 m/s^2 per minute of gap. Instead of reporting a single position, the model uses Monte Carlo sampling with 2000 draws to calculate the probability of entry as the fraction of sampled trajectories that intersect with the user defined zone. However, with the output being the result of a binomial distribution, it not only provides a probability estimate of zone trespassing but an uncertainty measure that depends on factors such as data quality, sea state, vessel dynamics or AIS gap duration, giving the analyst an assessment of the prediction reliability as well as the prediction itself.

Fig. 34. Predictive geofencing mode for the system, obtaining sea state data from AEMET sources.

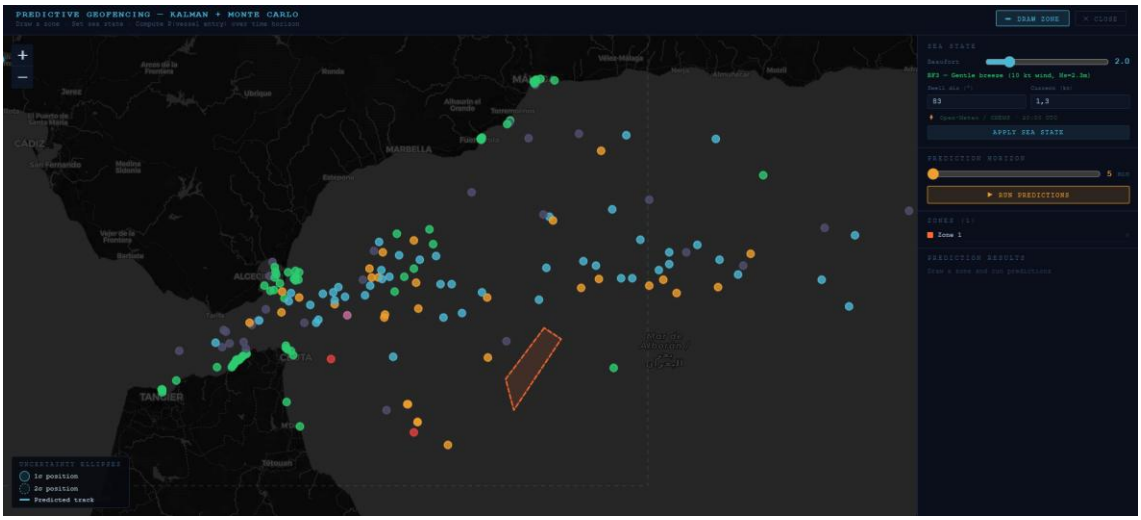

Fig. 35. Example polygon as established area of interest.

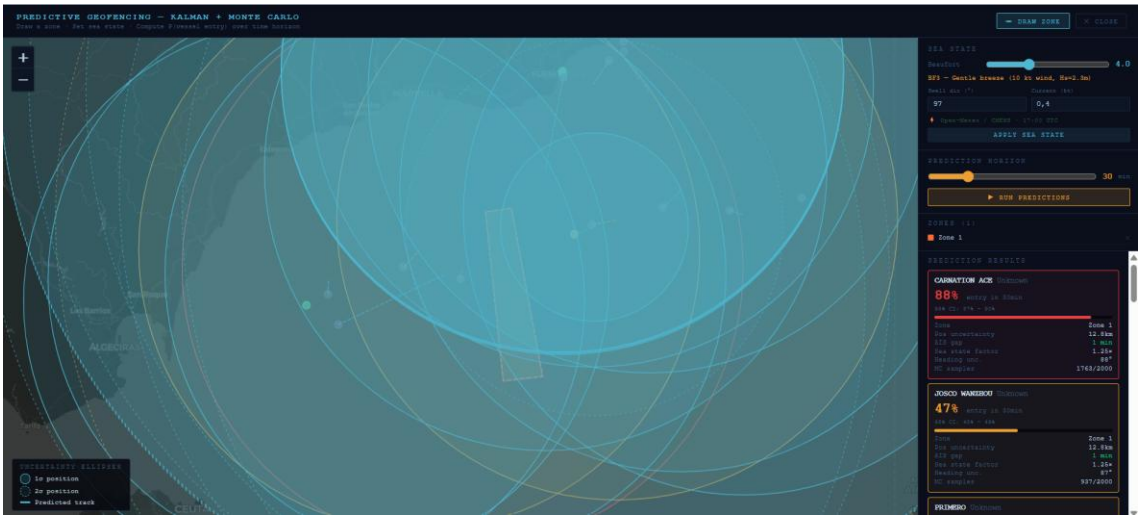

Fig. 36. Results of predictive geofencing analysis, showing vessels likely to intrude in the area of interest along the right of the screen.

4.4. Application to the Space Domain. Orbital Hybrid Intelligence Prototype.

To test whether the architecture developed in this chapter is specific to the maritime domain or is adequate for expansion into other areas of operation, a parallel prototype was developed applying the same philosophy to the space environment. The system mirrors the maritime model's structure, maintaining the same design principles of human oversight, explainability and reflecting uncertainty, replacing maritime specific inputs with their orbital equivalent. Vessel tracks becoming orbital parameters, AIS message becoming Two Line

Element (TLE) sets (the standard encoding for orbital parameters as defined by NORAD – North American Aerospace Defense Command), and flag or sanction risks being replaced by state affiliation or watchlist registers. The resulting system is built on the same layered JDL framework, with the same classifier structure and a clear separation of threat level from confidence scores.

The orbital prototype objective is the detection, and assessment of anomalous satellite behaviour, providing context for anomalous satellite behaviour, more specifically focusing on proximity operations, abnormal manoeuvring and sustained orbits near strategic assets that may be indicative of surveillance, harassment or preparation for hostile activity. Its intended users in an end state also coincide with the maritime model, including national security analysts, organisations with interests in the operational space with high value assets, commercial or government or Space Domain Awareness (SDA) organisations. As with the maritime system, it is important to reinforce that the detection of anomalous behaviour is not intended to be treated as irrefutable proof of hostile intent, but rather to trigger deeper analysis.

4.4.1.1. Data Fusion. JDL 0.

Where the maritime model ingests live AIS messages, the orbital model ingests real time TLE data fetched from providers such as CelesTrak GP API[60] (updated from real time data published by the 18th Space Defence Squadron) and SpaceTrack[61]. A TLE does not directly provide the current position of a satellite but instead contains a simplified description of a satellite's orbit at a specific point in time. In order to estimate the satellite's present location, each TLE is propagated forwards using the Simplified General Perturbations 4 model (SGP4), the standard orbital prediction model, producing position and velocity vectors. These vectors are then converted into operationally useful parameters such as latitude, longitude, altitude, inclination and orbital orientation. Together, these parameters provide a similar representation of satellite state to the positional and

[60] CelesTrak: A New Way to Obtain GP Data
[61] Space-Track.org

kinematic information obtained from satellite systems obtained from AIS data in the maritime prototype.

Additionally, parameters like AIS transmission gaps in the maritime system are translated into TLE update rates, with long data refresh gaps potentially indicating routine tracking rates or manoeuvres designed to evade data transmission. Satellite affiliation, or what operating state, military or civilian organisation the entities belong to are established through the use of a verified NORAD catalogue, built from documented incidents.

This layer is critical for contextualising threats in orbital space, for example the Cosmos 2581-2583 rendezvous operations[62], and backed by CelesTrak downloads covering satellite constellations such as GPS, GLONASS, BeiDou or Galileo, or crewed stations such as the ISS. Unrecognised objects are marked as unknown elements rather than attempting to infer classification for interpretation clarity.

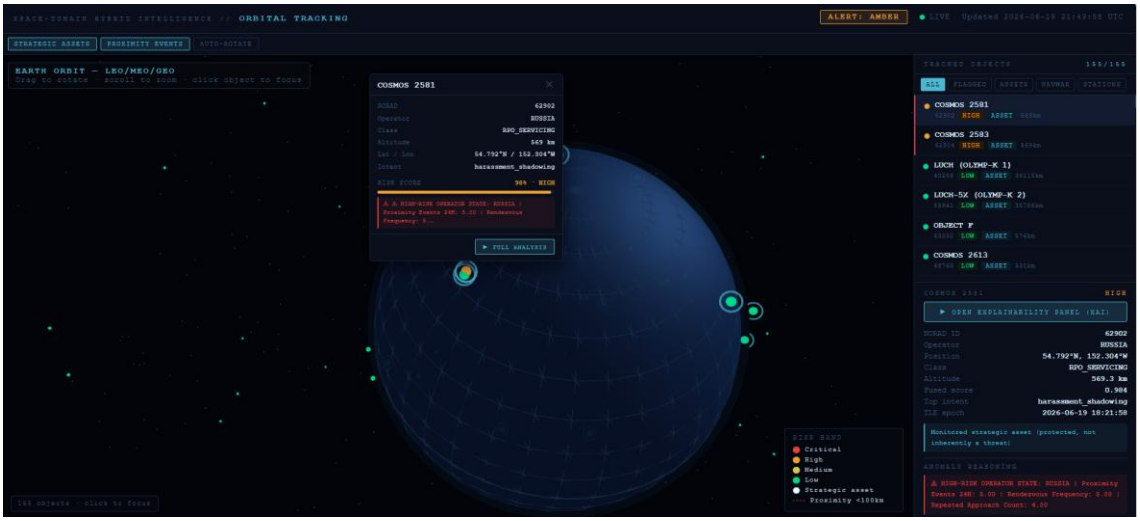

Fig. 37. COSMOS 2581 and COSMOS 2583 highlighted performing a rendezvous operation.

As illustrated in the tracking data in figure 35, COSMOS 2581-2583 satellites are seen performing proximity operations that are highlighted by

[62] SatTrackCam Leiden (b)log: RPO Galore: the complex proximity operations of Kosmos 2581, Kosmos 2582, Kosmos 2583 and 'Object F'

the model as such operations are clear examples of surveillance or hybrid activity training. This was corroborated using data from SATFLARE's tracking platform and imagery from SatTrackCam as shown below.

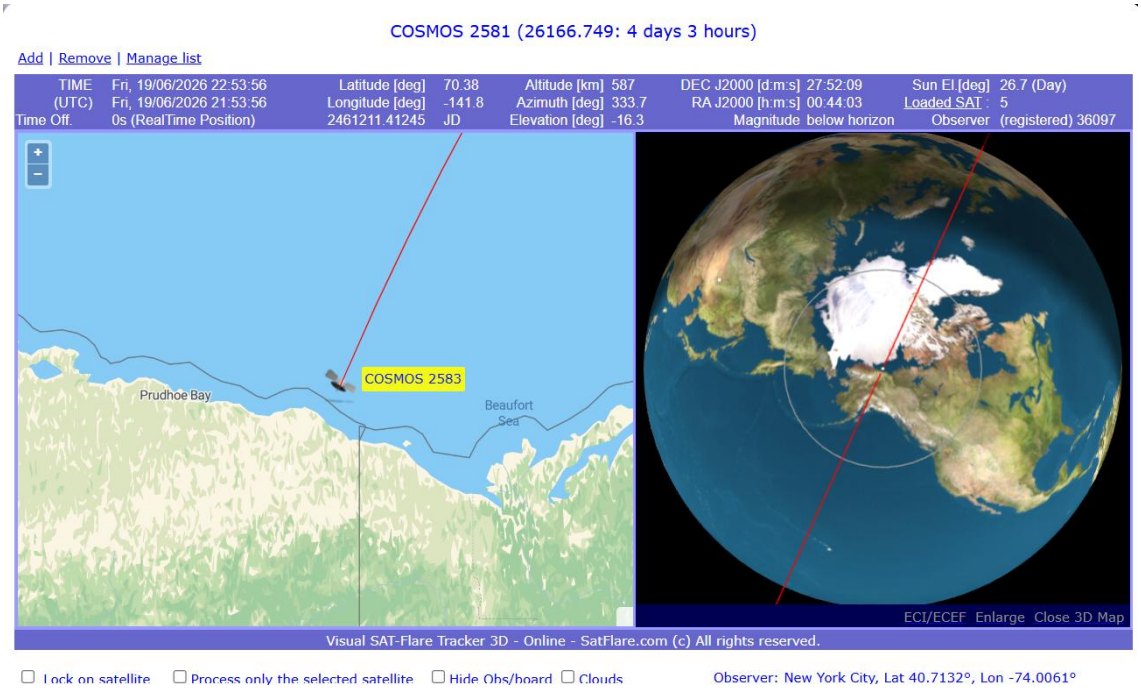

Fig. 38. SATFLARE data showing COSMOS 2581 and COSMOS 2583 engaging in close proximity operations as shown by the model.

Fig. 39. SatTrackCam image of COSMOS 2581, COSMOS 2582 and COSMOS 2583 performing close proximity operations.

The Space Domain Awareness system also incorporates functions that are not available for implementation in the maritime system, like direct manoeuvre detection. Because TLE data encodes average motion, comparing the orbital energy between two sets of data reveals whether a manoeuvre has been executed between observations. In order to do so, the semi major axis must first be derived using Kepler's third law.

First the mean motion given by the TLE, or the average angular rate at which a satellite travels around Earth in revolutions per day, must be converted to radians/second, the units required by Kepler's formula.

$$n_{rad} = n\left(\frac{2\pi}{86400}\right)$$

$$T^2 \propto a^3$$

$$a = \left(\frac{\mu}{n_{rad}^2}\right)^{\frac{1}{3}}$$

With T being the orbital period, n being the mean motion (the average angular rate at which a satellite travels around Earth in revolutions per day), a being the semi major axis of the satellite's orbit and μ being Earth's gravitational parameter 398600.4418 $\frac{km^3}{s^2}$.

Assuming a circular orbit, the orbital velocity is then calculated by applying the following formula.

$$v = \sqrt{\left(\frac{\mu}{a}\right)}$$

Producing an estimated orbital speed in km/s which can then be subtracted for consecutive TLEs, converting the result into metres per second before analysing.

$$\Delta v_{m/s} = 1000|v_2 - v_1|$$

Although simplified, this metric provides a physical backed indicator of orbital manoeuvres allowing the system to flag unusual behaviour quickly and without using unnecessary amounts of compute, much like the detection in course and vessel velocity for the maritime domain.

Proximity between tracked objects is also computed by bringing every tracked object to the same instant in time and measuring the straight line distance between their positions, with any pair closer than 100km being flagged and ranked depending on the real distance between them. To further clarify the nature of the approach, the relative closing speed is also reported, allowing the system to distinguish between a slow, sustained approach from a fast flyby.

4.4.1.2. Feature Extraction. JDL 1.

The information gathered is assembled into a 13 dimensional feature vector following the same model as the maritime system, except for it being built from orbital indicators instead. The vector includes the following indicators.

1. Time since last TLE update.
2. ΔV variance between TLE updates.
3. Number of orbital plane changes.
4. Count of close-approach events during a rolling 24 hour window.
5. Frequency of repeated approaches to the same target.
6. Affiliation risk based on operating state.
7. Deviation from expected behaviour.
8. Dual use or military classification risk depending on object type.
9. Unexpected manoeuvre flag.
10. Correlation with reports of electronic warfare or jamming.
11. Count of repeated encounters with the same target class.
12. Tracking gap before manoeuvre flag.
13. Interaction with watch listed strategic assets.

As with the maritime vector, missing contextual values are assigned neutral values instead of penalising ones, avoiding overstating suspicion of objects with fewer public information, leading to an overwhelming amount of false positives.

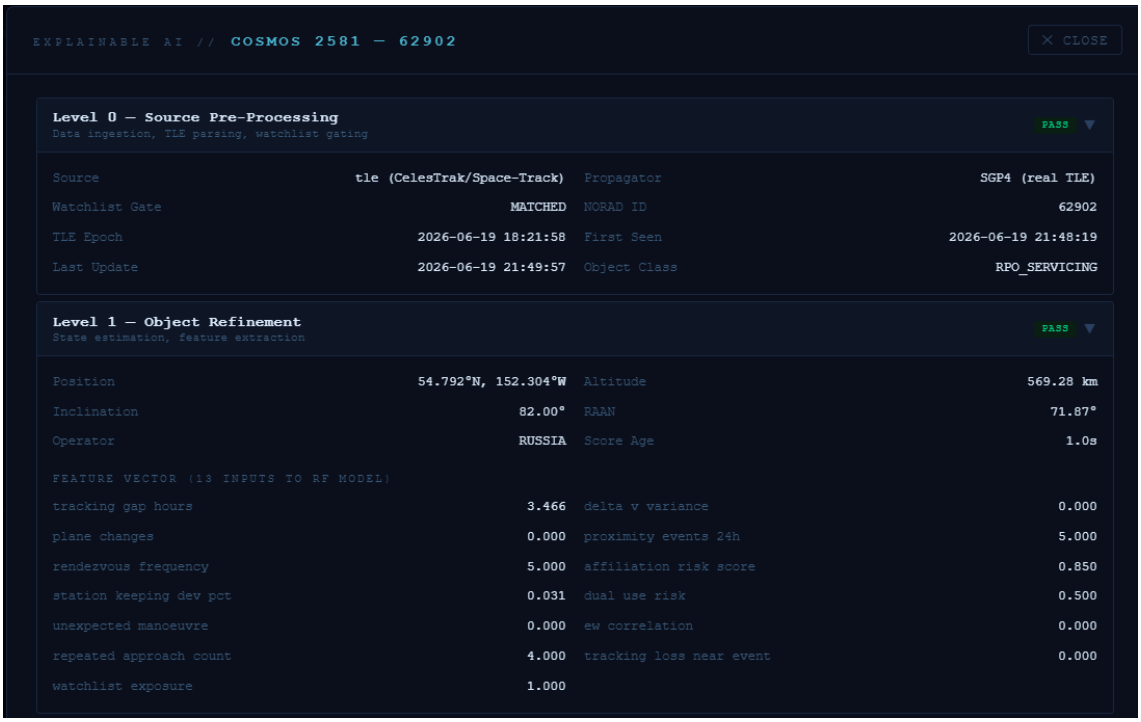

Fig. 40. Level 0 and Level 1 for COSMOS 2581 following the JDL framework.

4.4.1.3. Anomaly Scoring. JDL 2.

The orbital model reuses the same three method classifier which remains mostly unchanged. A Random Forest algorithm remains the primary supervised classifier, an Isolation Forest algorithm provides unsupervised outlier detection, especially useful for flagging anomalous behaviour that isn't included in training data, and an XGBoost classifier acts as a "second opinion", reducing the likelihood of bias. Due to the emerging nature of the operations, there is a lack of publicly available datasets for model training and validation, meaning that a synthetic dataset is generated specifically for this problem. However, it mirrors the SIAP dataset and follows the same philosophy, separating benign from hybrid behaviour and including plausible orbital behaviour such as a benign satellite exhibiting close approaches on a lower frequency than potentially suspicious objects (on the timescale of once every two weeks vs several per day for example). For clarity and consistency with established data

generation models, different probability distributions were assigned based on the data type generated. Count-based events such as plane changes, proximity events or repeated approach counts were modelled using a Poisson distribution[63], the conventional approach for discrete event counts. Time and magnitude-based features such as tracking gaps and Δv variance were modelled using an exponential distribution, which is commonly used for quantities that are usually small but occasionally have larger values[64]. This is adequate for the distribution of TLE updates, which are usually fairly frequent and involve minor manoeuvres, with delayed updates and larger orbital adjustments occurring less frequently. Bounded risk scores such as affiliation risk or dual use risk are modelled using a Beta distribution[65] since they are confined to a [0,1] interval, with benign scores being on the lower end of the scale and threatening scores being on the other. Lastly, binary indicators such as flags for loss of tracking before significant manoeuvres are modelled using a Bernoulli distribution as they are based on a yes/no system. As a result, the synthetic dataset represents a structurally sound model of how each feature type should behave, even though the exact rate and occurrences remain assumptions rather than empirically derived values.

The three scores are then fused using a weighted sum like in the maritime model:

$$S = 0.55\ P_{RF} + 0.20\ P_{XGB} + 0.25\ P_{IF}$$

The weighting chosen is slightly different than the maritime model's 0.5/0.3/0.2 split due to the absence of real labelled incident data for orbital hybrid activity, cause for the greater reliance on the unsupervised Isolation Forest. A watchlist is also used to mirror the sanction list lookup of the

[63] Cameron, A., & Trivedi, P. (1999). Regression analysis of count data. 2nd ed. In *Technometrics* (Vol. 41). https://doi.org/10.1017/CBO9780511814365

[64] Ross, S. M. (2014). *Introduction to probability models* (Eleventh edition). Elsevier.

[65] Ferrari, S., & Cribari-Neto, F. (2004). Beta Regression for Modelling Rates and Proportions. *Journal of Applied Statistics*, *31*(7), 799–815. https://doi.org/10.1080/0266476042000214501

maritime platform, with any object detected in sustained proximity to a watchlisted asset being assigned a fused score of at least 0.8, coinciding with the critical risk threshold. Model disagreement between supervised RF and XGBoost models is also flagged above a 0.25 margin, indicating operational ambiguity.

Confidence is also reported separately from threat scores, following the same principle as the previous model, a high score derived from outdated TLE data should not be interpreted in the same way as the same score derived from recently updated TLEs. In order to establish confidence levels, four parameters reflecting data quality are averaged to provide a value. It is worth noting that these parameters are averaged rather than weighted as done previously since no relevant studies have been produced indicating the dominance of one parameter over the others, being highlighted as an area for future development and potential model refinement once real incident data becomes available. The averaged components are:

1. Tracking source quality (Information derived from radar, optical methods or TLE only)
2. TLE data recency.
3. Model agreement (RF and XGBoost value output similarity)
4. Track history (scaled by the number of observations of that object prior to the moment of analysis)

To conserve the philosophy of explainable AI, every threat assessment is still accompanied by the SHAP values for the feature vector, analysing the contribution of each individual parameter to the risk assessment and its evolution.

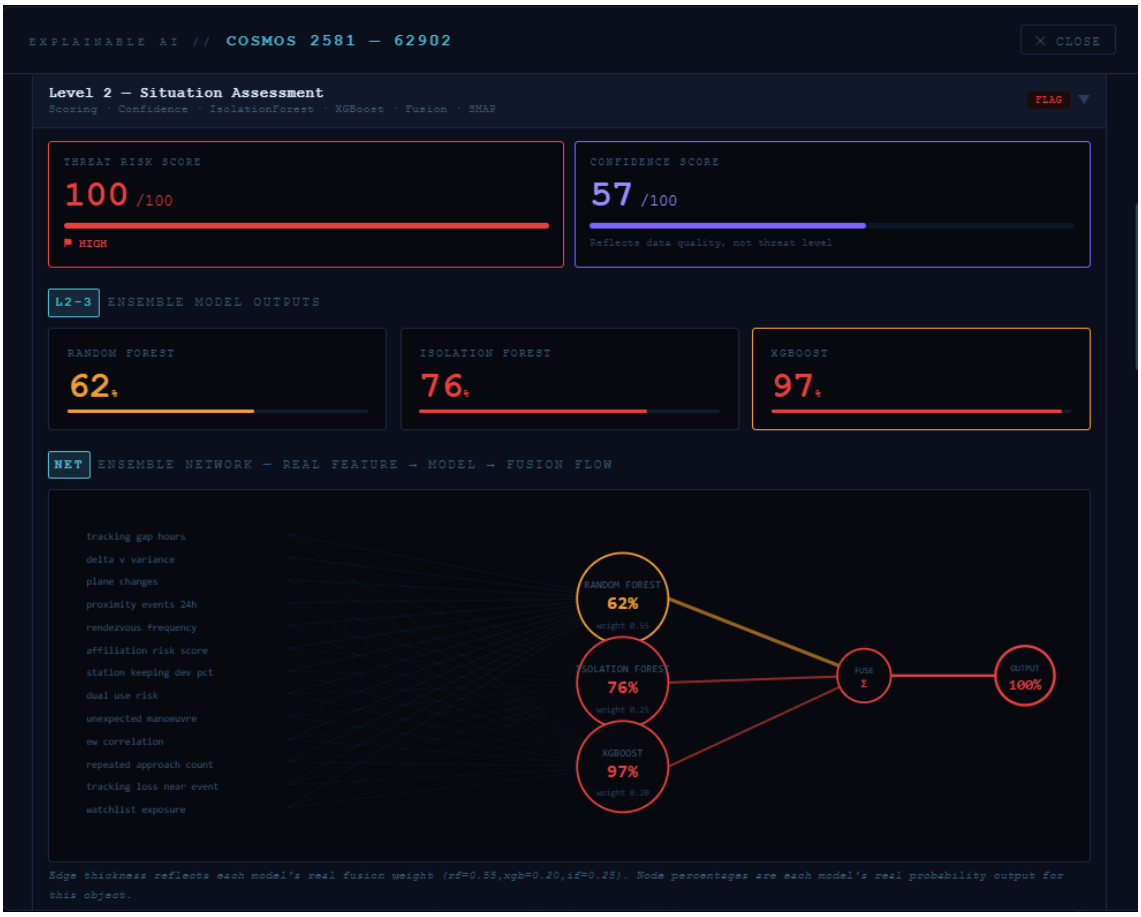

Fig. 41. JDL level 2 results for the initial analysis of COSMOS 2581.

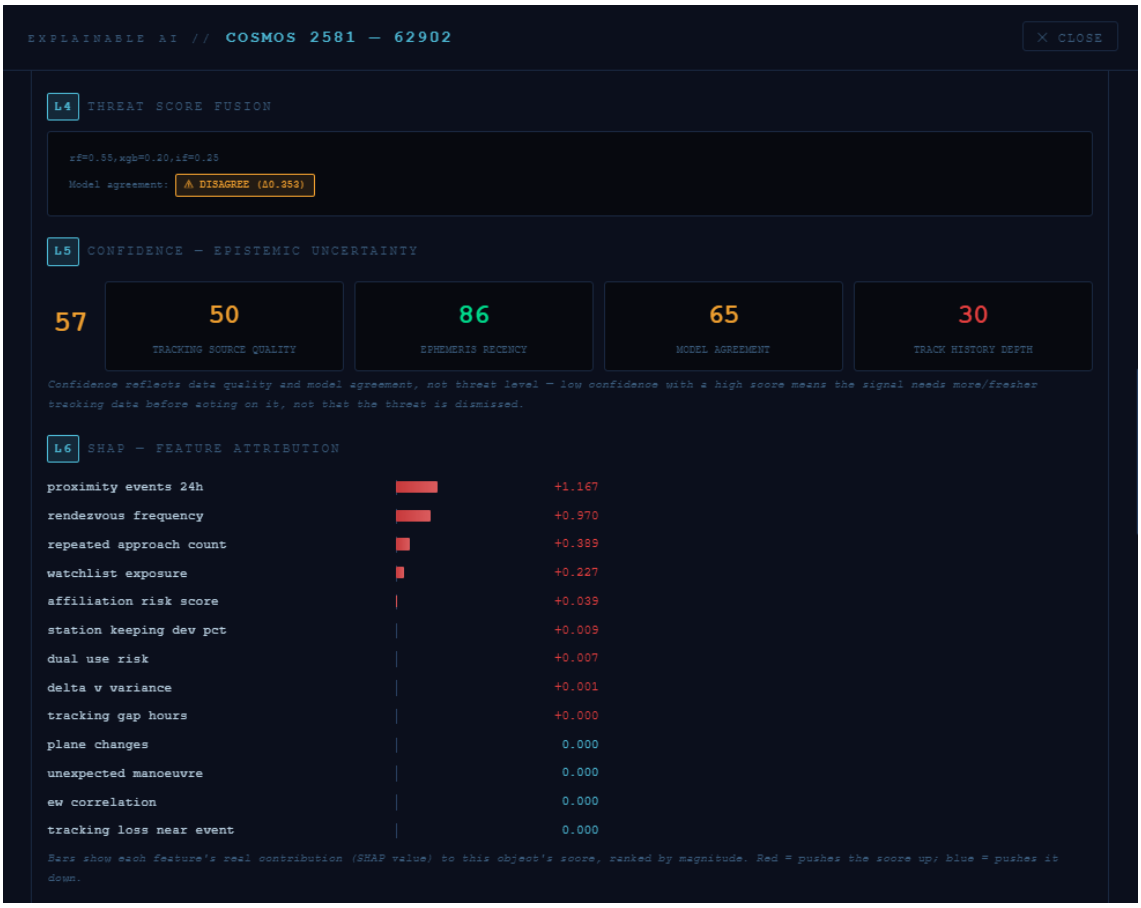

Fig. 42. SHAP values and confidence metric decomposition for COSMOS 2581.

Fig. 43. Feature importance decomposition by model.

4.4.1.4. Behavioural analysis and Intent Classification.

The orbital model maintains a rolling profile for every object, scored across six weighted behavioural indicators over a seven day window. This window is longer than the maritime 24 hour window due to the longer timeframes of space operations (at the moment) and the lower update frequency of TLEs compared to AIS. The tracked indicators are:

1. Sustained station keeping near strategic assets (such as the International Space Station)
2. Tracking gaps immediately followed by manoeuvres.
3. Inconsistent orbital plane changes.
4. Proximity operations and repeated approaches.
5. Sustained co-orbital positioning in regards with other assets.

6. Unexpected manoeuvre magnitude.

These indicators are fed into a rule based intent classifier (choosing rule based over learned for explainability reasons like in the maritime prototype) that assigns probabilities across six hypotheses.

1. Benign.
2. Surveillance.
3. Harassment or shadowing.
4. Pre-positioning.
5. Co-orbital ASAT preparation.
6. Electronic warfare support

Confidence in the top resulting hypothesis is reported alongside the margin between the two highest scoring hypotheses, with margins under 0.15 being flagged as ambiguous.

One other domain specific refinement applied to the model is the co-orbital ASAT preparation hypothesis only being valid when the nearest strategic asset belongs to a different operator than the approaching object, since the sustained proximity between two assets belonging to the same state is more likely to indicate servicing or inspection procedures.

Another area marked for future development is the development of the predictive geofencing framework for satellites, projecting the vessel's position forwards in time using Monte Carlo sampling to estimate the probability of entering protected zones like in the maritime system. Although all tracked objects are propagated to a common point in time using the SGP4 framework, no mechanisms have been implemented to forecast position or propagate uncertainty through time. However, if it was implemented, the forecasting accuracy would be a function of atmospheric drag, perturbation modelling and the degradation of TLE accuracy over time.

By analysing the correlation between inputs, variables and outputs of both models mentioned above, it is clear that the hypothesis of applying the same framework to operations of the same type across different domains is upheld. Every layer representing an object's state and history, as well as feature extraction, fusion, confidence and explainability remains the same with only slight adjustments for maintained accuracy. For instance, the dimensionality of the feature vector, the classifier architecture, fusion logic and explainability mechanisms remain the same across both systems. Layers diverging from each other did due to the structure of the examined area of operations, not because a generalisation problem, for example, with orbital tracking having lower update frequencies and therefore needing a longer behavioural window. Other instances of system modification include the averaging of the terms in the confidence formula rather than the weighing shown in the maritime system due to the lack of evidence showing the dominance of any parameter or the geofencing layer not being transferred due to the different kind of uncertainty required for orbital position propagation, even if the fundamental workings remain the same. This suggests that that the architecture's reasoning and representation can be generalised across domains, and that future extensions of this framework to other areas of operation should be developed in the same manner, by first transferring the overall structure and only then modifying domain specific parameters where the layered stages depend on the physics of the observed environment.

Fig. 44. Threat assesment and behavioural pattern analysis for COSMOS 2581.

4.5.System Validation.

Since both systems operate on synthetic training data due to the absence of publicly available labelled hybrid threat datasets (modelled using the SIAP dataset in the case of the maritime prototype), a validation framework was developed to determine the accuracy of the scoring algorithm. In order to calculate the metrics describing the system, a sample of 1000 vessels and satellites were taken for each model across three classes, “normal”, “suspicious” and “threat”. In both cases, the dataset is divided into a 75% training set and a 25% held-out test set, which acts as a reserved dataset that the model doesn’t have access to during training and serves to evaluate the model performance as a “final test”. It is important to note

that the data was deliberately stratified to ensure a similar distribution in the training and held out sets to ensure fairness in the testing. The maritime framework utilises a skewed 60/25/15 normal, suspicious, threat distribution to represent crowded, high traffic commercial shipping lanes flooded with everyday traffic. However, the space dataset includes a higher proportion of "suspicious" and "threat" behaviours following a 50/30/20 distribution, since the baseline environment is stricter, meaning any slight deviation is statistically significant.

4.5.1. **Interpretability of Evaluation Metrics.**

Before presenting the results, it is worth explaining what the evaluation metrics represent. Once the data is introduced into the system, the classifier sorts the inputs into the three categories mentioned in section 3.5. Following the classification, the sorting is evaluated to determine the precision of the classification, or in other words, how many items have been correctly assigned their corresponding labels, with a precision of 1.0 signifying all of the inputs for that category have been correctly assigned to their respective labels. The second metric is recall, or the fraction of real threats that the system actually identifies. In this case, a recall of 1.0 means that every threat has been successfully identified. Both of these values are measured using the F-1 score, a metric that penalises systems that sacrifice one to inflate the other. Continuing with this example, a system that flags every vessel as a threat would achieve a high recall score but very poor precision.

In order to map the values across all three classes, the confusion matrix is built and studied. Each row on the confusion matrix represents what the vessel actually is, with each column representing what the system predicted. A perfect system would show all of the numbers on the diagonal and zeros everywhere else, as off diagonal entries reflect classification errors. The last system metric is the macro F1, that averages the F1 score across all three classes equally, regardless of the number of items per class, an important condition to note since overall accuracy would be biased towards classifiers that ignore the minority class (a system that labelled

every vessel as normal would score 60% accuracy on this dataset while being inaccurate). Finally, a baseline is created by obtaining performance values from a classifier that ignores all features and always predicts the most common class, providing a performance floor below no useful system should fall, establishing a metric for model improvement.

4.5.2. **Maritime System Results**

The full three model pipeline achieved a macro F1 value of 0.931 against a baseline value of 0.250, representing a 0.681 improvement over a classifier with no predictive capability. However, when breaking the F1 metrics by class, “normal” achieves a precision of 1.000 and a recall of 0.980, “suspicious” achieves a precision of 1.000 and a recall of 0.873 and “threat” achieves a precision of 0.771 and a recall of 1.000.

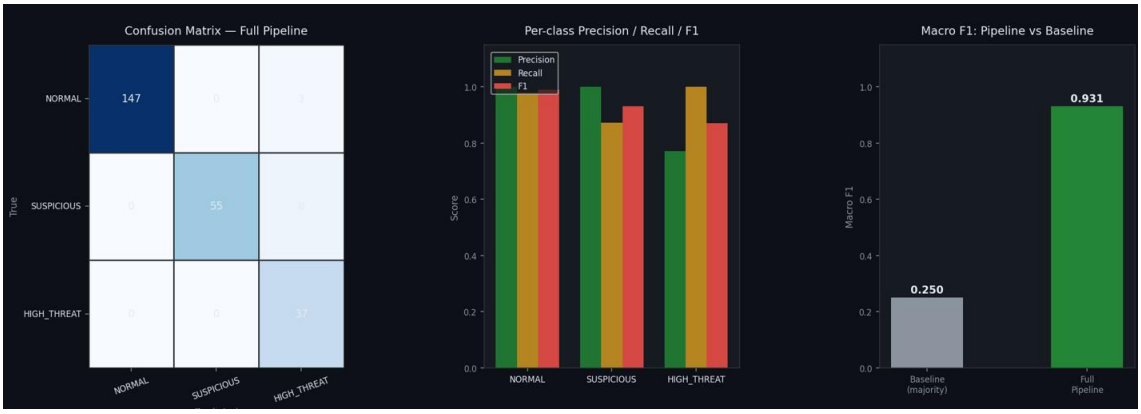

Fig. 45. Confusion matrix, precision, recall and F1 values by class and macro F1 value compared to the baseline for the maritime system.

The logic behind the class results is significant, mattering more than global metrics. The confusion matrix in figure 43 shows that the system flagged 48 vessels as high threat in the test set, of which, 37 were genuine, 8 were suspicious vessels that had been escalated one tier and 3 were normal vessels that had been escalated 2 tiers. By examining these results, it can be concluded that the precision reflects the system being conservative by design, accepting a small rate of upward classed false alarms to ensure that

no genuinely threatening vessels are undetected. The recall of 1.000 confirms this, as all 37 truly threatening vessels in the test set were correctly identified. From an intelligence support perspective, this validates the system reasoning, as the cost of a type II error (missing a threat) can be catastrophic, whereas the cost of a false alarm or type I error can be acceptable in comparison.

4.5.3. **Space Domain Results**

The space system model achieved a macro F1 of 0.904 against a baseline of 0.222, a 0.682 improvement, achieving very similar results to the maritime system. The threat class achieved a precision of 0.725 and a recall value of 1.000, again demonstrating a conservative approach intended to detect every genuine hybrid threat.

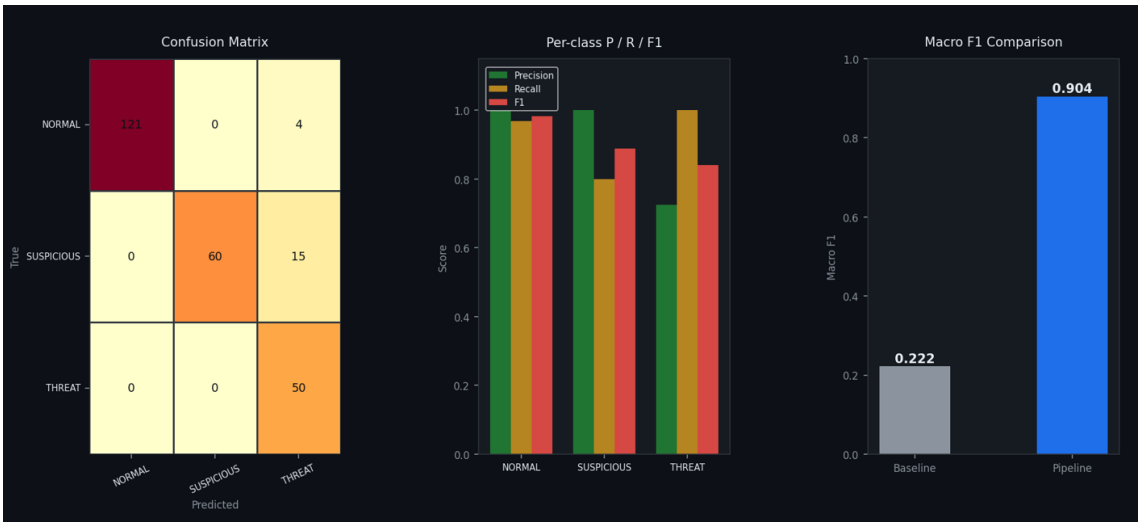

Fig. 46. Confusion matrix, precision, recall and F1 values by class and macro F1 value compared to the baseline for the space analysis system.

This precision degradation can be seen clearly in the confusion matrix shown in figure 46, with 4 normal and 15 suspicious satellites being escalated to “threat” levels. This scoring is also consistent with the hard floor of 0.8 for any object shadowing watch listed assets, meaning that any

normal satellites coincidentally following watchlisted orbits may be flagged as threatening, driving down precision.

4.5.4. Robustness and Sensitivity Analysis.

Another model parameter worth examining is the system's weight and threshold values, which were assigned through structured reasoning. Firstly, the Isolation Forest algorithm received the lowest weight due to its unsupervised nature and susceptibility to false positives, with the Random Forest algorithm receiving the highest weight as the primary supervised classifier. The 0.8 floor for sanctioned of manually flagged objects was chosen to be sufficiently high to ensure that deliberately manufactured positive indicators would not mask a vessel appearing on a sanctions list. However, in order to validate the weights a sensitivity analysis was ran across both systems, varying each fusion weight by ± 0.100 and each threshold by ± 0.050 while holding all other parameters fixed in place. As seen in figures 45 and 46, weight perturbation changes showed no shift across F1 values for both systems and all tested configurations, demonstrating that the system is stable and robust for the chosen weights. By examining the effect of the 0.8 floor sensitivity value on the macro F1 as seen in figure 47, it can be observed that the resulting F1 value also remains constant for changes in value in an interval of [0.7, 0.9], proving to be stable. Similarly, threshold value changes for the space-focused system show no change to F1 values, with only a minimal change of 0.0038 being triggered when the normal/suspicious boundary or both the normal/suspicious and the suspicious/threat boundaries are lowered by 0.05 for the maritime system. This negligible variance confirms that the model's performance does not depend on hypersensitive tuning, and that instead, the model parameters are situated within a broad "plateau" for which classification is stable.

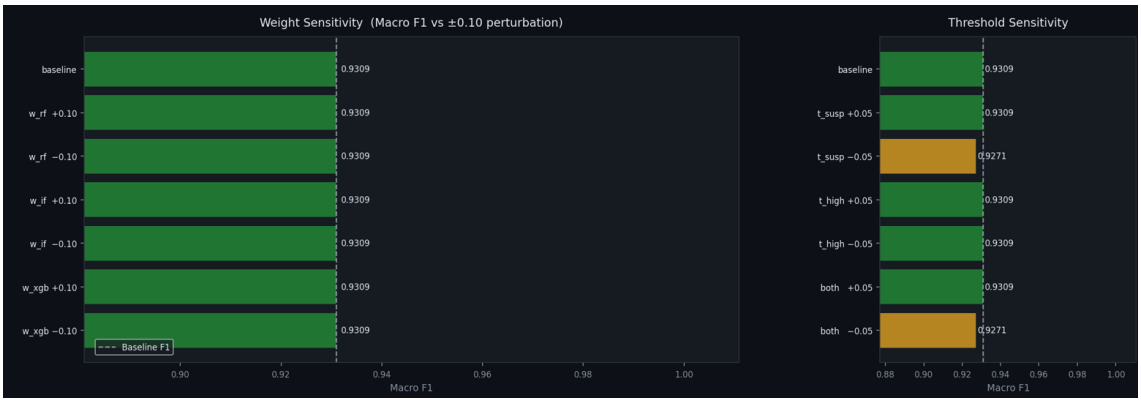

Fig. 47. Weight and threshold sensitivity analysis for the maritime system.

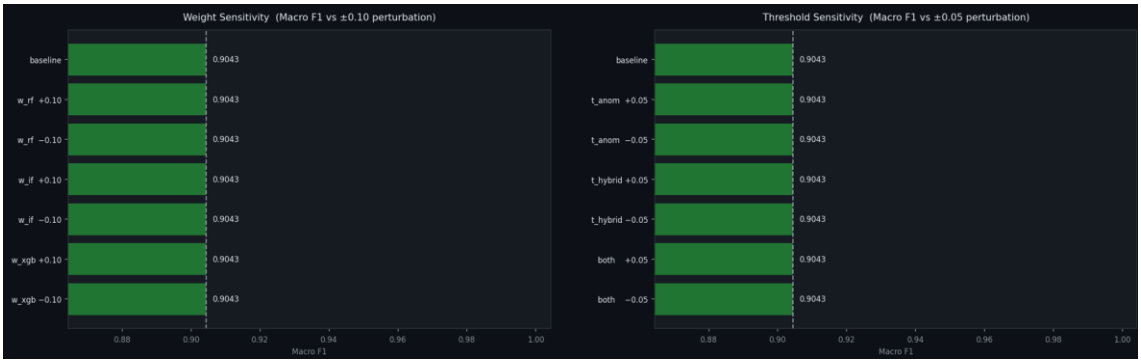

Fig. 48. Weight and threshold sensitivity analysis for the space-focused system.

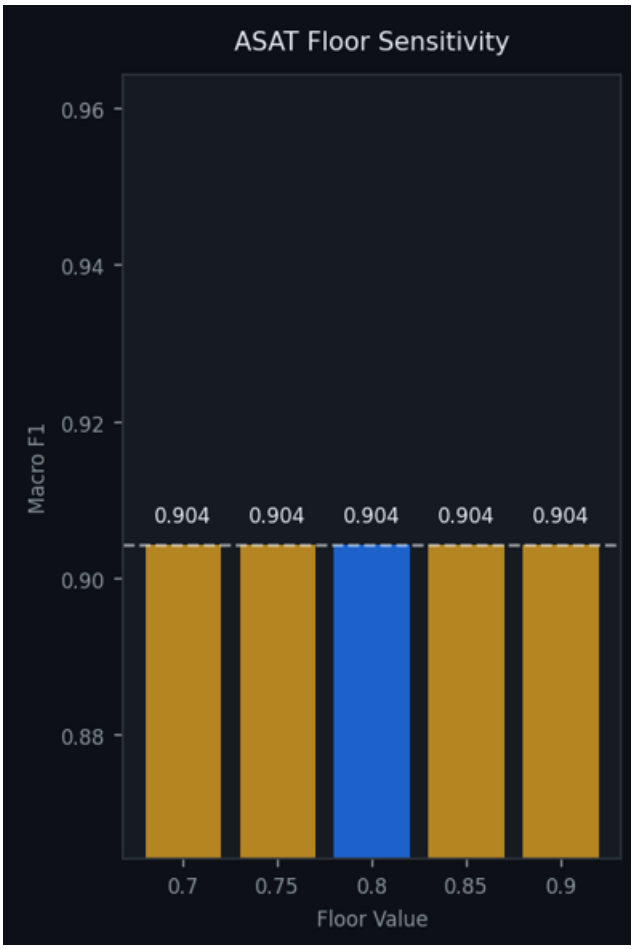

Fig. 49. Floor value sensitivity analysis.

4.5.5. Summary of Results.

The validation results across both models confirm the three main claims made at the beginning of the system design. First, the pipeline successfully distinguishes between behavioural classes that current monitoring systems struggle to separate. The suspicious tier, which corresponds to ambiguous grey zone activity characteristic of hybrid operations is correctly identified in the vast majority of cases across both systems, demonstrating that the relevant behaviour at an intermediate level is detectable from open source data. Second, the system's conservative architecture, designed to prioritise high recall over precision at the "threat" classification produces zero missed detections for threatening objects across both held out test sets.

Finally, the sensitivity analysis confirms that the hand-tuned parameters do not introduce fragility into the model, with performance being stable across all tested weight shifts.

However, the synthetic nature of the training data, the controlled conditions of the validation environment and the absence of adversarially manipulated inputs introduce opportunities for system development to improve model resilience and impose limits on what could be generalised to a live fully functioning system. Additionally, because the current testing baseline does not simulate intentional manipulation and attacks against the models themselves, the system's vulnerability to targeted poisoning or subversion remains unquantified. Nevertheless, the results do prove that the architectural decisions explained in this chapter are coherent, backed by evidenced performance claims and providing a proof of concept for AI-assisted hybrid threat detection. The limitations of this approach and the development pathways that follow from them are further addressed in Chapter 5 below.

Chapter 5. CONCLUSION.

4.1 Assessment Against Project Objectives.

This project set out to demonstrate that AI applied to open-source maritime hybrid data could support intelligence analysts in detecting and anticipating hybrid threats, shifting the analytical posture from responding to incidents as they happen in a reactive manner, to a proactive mindset. When evaluated against the project objectives, the prototype provides evidence that this goal is achievable both technically and conceptually. The evolution from incidental anomaly detection towards intent inference through behavioural analysis has been directly addressed through the development of the system with the combination of unsupervised anomaly detection through Isolation Forest algorithms and supervised classification algorithms such as Random Forest or XGBoost, as well as temporal profiling and rule-based intent classification. Together, these techniques provide assessments that are both grounded and operationally useful, more than if they had been applied on their own.

The Strait of Gibraltar case analysis also confirmed the hypothesis of composite scoring being relevant in a complex environment, especially focused on the co-occurrence of behavioural indicators. The selection of the three machine learning classifiers is also particularly significant. The Isolation Forest algorithm detected vessels exhibiting anomalous data outputs relative to normal traffic and did not appear in the Random Forest's training distribution, which aims to target the adaptable nature of hybrid operations. Another characteristic that serves as a distinct advantage over current systems is the RF vs XGBoost disagreement measurement, helping to identify cases where the evidence trace was genuinely ambiguous, turning ambiguity and uncertainty into an analysis tool and indicator in a way that single classifier systems would not be able to.

The project also demonstrated the viability of the use of open-source data hypothesis, using AIS feeds, ITU vessel registries, and public sanctions databases

and port state control records to populate the 13 dimensional feature vector used for risk scoring without having access to classified intelligence. While not particularly relevant to state agencies, this is directly relevant for the broader system user base, including organisations where access to classified feeds is restricted or operating under information sharing constraints. The focus on explainability, within the domain of eXplainable AI, also guarantees feasibility of adoption across organisations as desired, as well as an easier integration into existing workflows. As opposed to other standalone anomaly detection tools, designed to replace analysts, explainability is one of the main design constraints of the system. The SHAP efficiency axiom, ensuring that contributions of each feature to the model are balanced, without any one feature disproportionately influencing the outcome, guarantees fairness and removes biases, providing analysts with an exact account of the model's prediction instead of an approximation. Combined with the confidence score being explicitly separated from the threat level analysis, and the intent classifier's transparency and temporal evidence chain, the system provides a higher degree of analytical transparency, one of the main requirements set out in the concept brief as it is a prerequisite for adoption in intelligence environments where assessments must be defensible.

Integrating uncertainty quantification through the use of Kalman filtering also aids explainability, offering grounded and transparent outputs as answers to analyst's questions. It is important to note that, being grounded on a strict probabilistic mathematical base, results obtained from the estimation process are transparent in communicating confidence assessment metrics rather than delivering opaque precision measurements, with the 95% Monte Carlo confidence interval presenting the analysts with a quantitatively grounded basis for deciding whether a predicted geofence breach warrants further investigation.

4.2 Recommendations to Policy Makers.

The findings of this project carry several implications for policy makers at the European and NATO level. The most significant is the case for embedding AI-assisted analysis within existing intelligence workflows within organisations such as MARCOM or FRONTEX. Existing intelligence collection structures face a

management problem, as the volume of data and evidence available exceeds the analytical capacity of human only workflows, and will only continue to worsen in the age of hyperconnectivity. The system developed in this paper and the analysis of the state of affairs outlines how this gap can be closed without replacing human analyst judgment with a human oriented, explainability first approach. Integrating these models into decision chains becomes less challenging when human oversight is built into the system, positioning the analyst at the end of the data pipeline, also helping with positioning the developed systems in the European AI regulatory debate. By retaining explainability as a first principle, integrating characteristics such as SHAP based attribution and confidence scoring also reduce the risk of overreliance on automated scoring, as they are designed to communicate uncertainty instead of concealing it. However, implementing this recommendation is not solely a question of technical readiness. Embedding AI assisted analysis into intelligence workflows requires testing and certification processes to ensure that analysis tools can be trusted within classified operational environments, a process that can take years to implement in defence institutions, accustomed to slower technology adoption cycles. Introducing AI systems into classified environments also requires absolute guarantees on uncompromised operational security, ensuring both software and hardware introduced into secure facilities meet the necessary requirements. Additionally, analysts trained in traditional tradecraft may be reluctant to incorporate machine-generated risk scores into their assessments, particularly in environments that reward caution over experimentation. Finally, coordinating access permissions and data sharing across agencies with different security clearances adds a layer of friction that can't be resolved solely on a technical basis. Overcoming these problems will require training and a unified organisational shift, developing standardised protocols that explicitly demonstrate model reliability to analysts, bridging the gap between traditional tradecraft and automated decision support systems.

A second policy recommendation concerns data standardisation and interoperability of systems across the NATO environment. The performance of the data analysis pipeline is directly correlated to the consistency and availability of inputs from AIS or other databases, with current fragmentation across European systems in AIS collection coverage, vessel database access or port state control

record publications representing a major capability gap. Information is often presented according to the operational needs of the organisation that collects it, creating interoperability challenges and forcing systems to parse and assemble data from a variety of different sources that may contain contradictions instead of facilitating a single point for data exports. Coordination to establish common data standards and processing frameworks (such as the JDL architecture), as well as facilitating cross border information sharing, much like the processes in place for communicating financial sanctions, would have a significant effect on improving the operational efficiency of systems in this area. However, this recommendation carries significant challenges and might not prove to be immediately actionable. According to the European Commission, member states have been hesitant to share raw intelligence for reasons of operational security rather than technical limitations[66]. Furthermore, commercial maritime data providers have limited incentives to standardise outputs, as that would entail losing differentiation from their competitors and disclosing their full system capabilities. NATO's track record of negotiating common technical standards also suggests that even when the operational benefit is clear and agreed upon, reaching consensus across the alliance can be a multi-year project rather than an immediate fix. For example, the approval of STANAG 4586 for the standardisation of interfaces for UAV control system interfaces to enable interoperability[67] was first brought forward as a concept in 1999, but was not published until 2003 as the first edition and did not fully align with modern cross-platform requirements until edition 3, published in 2012.

The passive deterrence dimension is also crucial in combating hybrid threats, especially as it is often relatively inexpensive compared to relative hard defence measures. Improved monitoring and documentation of suspicious behavioural patterns communicates to grey zone actors that the detection risk associated with covert maritime operations has increased, even in the absence of immediate enforcement. The credibility of the deterrence is based on the common knowledge

[66] (PDF) Maritime Surveillance and Information Sharing Systems for Better Situational Awareness on the European Maritime Domain: A Literature Review. (n.d.). In *ResearchGate*. https://doi.org/10.1007/978-3-030-65722-2_8

[67] STANAG 4586–Standard Interfaces of UCS for NATO UAV Interoperability

of a known capability, without disclosing specific features that could lead to adversarial adaptability. Therefore, policy frameworks that support the transparent communication of monitoring capabilities within the appropriate classification limits have an increased strategic value that acts as a multiplier to their direct operational utility.

A third recommendation refers to the governance framework for AI assisted intelligence recommendations. The human in the loop architecture upon which the system developed in this paper is built represents an operational necessity due to the intelligence context but could be expanded to AI development in other areas. Organisations working on or deploying AI assisted analysis tools should establish policies governing the role of automated recommendations within the decision making process, reaching a compromise in a way that does not stifle AI development. By integrating explainable AI as a fundamental value of AI systems and focusing development on human support functions, it is possible that regulation around system development could be relaxed, boosting sector growth and positioning Europe as a capable leader. Equally important is preserving analyst accountability. AI systems should support analytical judgement rather than displacing it, ensuring that intelligence conclusions remain attributable to human decision makers, especially at a policy and strategic levels.

Excessive reliance on AI integration into intelligence workflows also creates strategic implications in the long term, including cognitive dependency in which analysts become increasingly inclined to believe machine generated outputs without sufficient scrutiny. Over time, similar issues could lead to the erosion of critical analyst skills such as the capacity to challenge information received from sources, leading to biased assessments. Furthermore, this also opens the door to new types of cognitive warfare, tampering with model training or data weights (or even changing the way the models address analysts, for example, increasing sycophantic traits) to change the direction of an agencies intelligence analysis, or even hampering the ability to operate in case of model failure or disconnect. Transparency at the technical level must also be matched at the organisational level. Explainable AI can provide visibility into how confidence scores are generated, but without documentation standards or oversight structures, the

practical transparency effort may have limited value. Effective governance therefore requires the integration of technical transparency through techniques such as the integration of SHAP values with analyst training and unified frameworks capable of supporting trustworthy AI driven analysis.

4.3. Limitations and Future Development.

The developed system has several limitations that future development should address, with the most significant one being the training of the Random Forest and XGBoost (both supervised models) on synthetic data. While the training data parameters were derived and calibrated from the published SIAP paper, the absence of labelled incident data means that classifiers have not been trained against confirmed hybrid activity. This limitation only applies to the orbital classifier model training process and not to the data it scores while running, which is an important distinction. However, while the scarcity of labelled incidents is a characteristic problem for the hybrid operations domain, the current system addresses it by treating probability outputs as relative risk indicators rather than calibrated probabilities.

The data model relying on AIS is also a system limitation. Integration with Synthetic Aperture Radar data and commercial satellite data, as well as other data sources such as RF emission tracking would significantly reduce the false negative rate for vessels exploiting the limitations of AIS coverage. Furthermore, the ownership graph analysis with data obtained from ITU MARS records would also benefit from integration with commercial ownership databases to improve traceability with regards to shell companies among other situations. However, at the time of system development, few databases providing the required imagery or information matching the open-source criteria could be found.

Finally, the system currently lacks a mechanism for incorporating analyst feedback into updating the model, offering training opportunities derived from human judgement rather than artificial generation. This would be realised by

having analysts interact with the model by confirming, overriding or adding context to datapoints, improving the performance of the classification systems and adding extra adaptability to the model, serving as a continuous improvement mechanism. However, the main challenge holding this opportunity back is not technical rather than procedural, as creating the environment and data sharing mechanisms required for the system to function poses a challenge in itself, analogous to the creation of forecasting systems like the PolicyMAP structures mentioned in this paper.

Regardless of these limitations, the system proves that the project objectives established in the hypotheses and concept brief are achievable using current data processing systems and open sources. The findings also suggest that the main challenge in interpreting hybrid threats and anticipating grey zone operations is not a lack of data but rather, transforming the increasingly growing volume of data into actionable intelligence. As hybrid actors continue to exploit ambiguity, the ability to identify behavioural patterns, infer intent and anticipate future activity will become increasingly important for supporting decision making, especially as hybrid operations grow in their current applications and expand further into other environments such as space.

This paper also set to illustrate how artificial intelligence can be integrated into human workflows to support rather than replace human judgement and analysis, remaining transparent and explainable. At the same time, it highlights the strategic dependencies that will enable further AI development to contribute to national security, including access to computational infrastructure, energy security and technological sovereignty. Ultimately, the transition from reactive monitoring towards predictive maritime intelligence represents both technological evolution and a strategic shift. Organisations capable of identifying emerging grey zone threats before they are allowed to inflict damage will possess a significant operational advantage in increasingly contested and complex environments, making predictive intelligence a critical component of future defence and intelligence capabilities.

Chapter 6. BIBLIOGRAPHY.

> foia-home > foia-declassified-nro-programs-and-projects. (n.d.). Retrieved 14 April 2026, from https://www.nro.gov/foia-home/foia-declassified-nro-programs-and-projects/

A BRIEF HISTORY OF THE JOINT DIRECTORS OF LABORATORIES DATA FUSION GROUP - ProQuest. (n.d.). Retrieved 24 April 2026, from https://www.proquest.com/openview/9d14d9abdb72ff21b559591ea509bb8b/1?pq-origsite=gscholar&cbl=6699140

Airehenbuwa, B., Hasan, T., Sarkar, S., & Guin, U. (2025). *Advancing Security with Digital Twins: A Comprehensive Survey* (arXiv:2505.17310). arXiv. https://doi.org/10.48550/arXiv.2505.17310

AI's Impact on International and Geopolitical Dynamics. (n.d.). *Brookings*. Retrieved 20 May 2026, from https://www.brookings.edu/tags/ais-impact-on-international-and-geopolitical-dynamics/

AIS in a Historic Perspective. A History of the Identification of Ships. - Royal Institute of Navigation. (n.d.). Retrieved 25 March 2026, from https://rin.org.uk/blogpost/1706945/357655/AIS-in-a-Historic-Perspective-A-History-of-the-Identification-of-Ships

AIS spoofing: The fast track to sanctions | Kpler - Nov 10, 2025. (n.d.). Retrieved 29 March 2026, from https://www.kpler.com/zh-sg/blog/ais-spoofing-fast-track-to-sanctions

Akidau, T., Bradshaw, R., Chambers, C., Chernyak, S., Fernández-Moctezuma, R. J., Lax, R., McVeety, S., Mills, D., Perry, F., Schmidt, E., & Whittle, S. (2015). The dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. *Proceedings of the VLDB Endowment*, *8*(12), 1792–1803. https://doi.org/10.14778/2824032.2824076

Akıllı, E. (2026, April 15). *Cutting cables: How war in Red Sea could trigger digital collapse | Opinion*. Daily Sabah. https://www.dailysabah.com/opinion/op-ed/cutting-cables-how-war-in-red-sea-could-trigger-digital-collapse

Ali, I., Stewart, P., Ali, I., & Stewart, P. (2026, January 7). Exclusive: US seizes Venezuela-linked, Russian-flagged oil tanker after weeks-long pursuit. *Reuters*. https://www.reuters.com/business/energy/us-seizing-venezuela-linked-oil-tanker-after-weeks-long-pursuit-2026-01-07/

All Together Now: China's Militia in 2025 | Asia Maritime Transparency Initiative. (n.d.). Retrieved 6 March 2026, from https://amti.csis.org/all-together-now-chinas-militia-in-2025/

Amazon Wins Back NSA's $10B Hybrid Cloud Deal—Potomac Officers Club. (2022, April 29). https://www.potomacofficersclub.com/news/amazon-wins-back-nsas-10b-hybrid-cloud-deal/

Anh, D. T. (n.d.). *The Jevons Paradox in AI Infrastructure: Efficiency, Rebound, and the Limits of Cost Reduction*.

Artificial Intelligence and the Future of Strategic Stability. (2026, March 9). *Texas National Security Review*. https://tnsr.org/roundtable/artificial-intelligence-and-the-future-of-strategic-stability/

Attard, F. G. (2020). *The Duty of the Shipmaster to Render Assistance at Sea under International Law*. Brill | Nijhoff. https://doi.org/10.1163/9789004438255

Automatic Identification System (AIS). (n.d.). *Great Lakes St. Lawrence Seaway System*. Retrieved 25 March 2026, from https://greatlakes-seaway.com/en/commercial-shipping/transiting-the-seaway/automatic-identification-system-ais/

Balduzzi, M., Pasta, A., & Wilhoit, K. (2014). A security evaluation of AIS automated identification system. *Proceedings of the 30th Annual Computer Security Applications Conference*, 436–445. https://doi.org/10.1145/2664243.2664257

Barberi, E., Chillemi, M., Cucinotta, F., Raffaele, M., Salmeri, F., & Sfravara, F. (2025). Leveraging Artificial Intelligence for Real-Time Risk Detection in Ship Navigation. *Applied Sciences*, *15*(21), 11674. https://doi.org/10.3390/app152111674

Bar-Shalom, Y., X., R. L., & Kirubarajan, T. (2001). *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc.

Beneath NATO's Radars: Unaddressed Threats to Subsea Cables | Strategic Technologies Blog | CSIS. (n.d.). Retrieved 14 May 2026, from https://www.csis.org/blogs/strategic-technologies-blog/beneath-natos-radars-unaddressed-threats-subsea-cables

Berbić, I., Stupalo, V., Kavran, N., & Bukljaš, M. (2023). Overview of machine learning methods in maritime traffic monitoring based on Automatic Identification System. *Transportation Research Procedia*, *73*, 220–226. https://doi.org/10.1016/j.trpro.2023.11.911

Bowman, C. L., White, F. E., Blasch, E. P., & Llinas, J. (n.d.). *Independent Consultant Alexandria, VA, USA alaneilsteinberg@gmail.com*.

Buncic, D. (2016). Superforecasting: The Art and Science of Prediction. By Philip Tetlock and Dan Gardner. *Risks*, *4*, 24. https://doi.org/10.3390/risks4030024

C2coe, N. (n.d.). *Towards Improved Pathways for Human–Machine Teaming in Command and Control at the Operational Level – NATO C2COE*. Retrieved 20 May 2026, from https://c2coe.org/download/towards-improved-pathways-for-human-machine-teaming-in-command-and-control-at-the-operational-level/

Cameron, A., & Trivedi, P. (1999). Regression analysis of count data. 2nd ed. In *Technometrics* (Vol. 41). https://doi.org/10.1017/CBO9780511814365

Can Maritime Law inform Space Law in addressing international commons? (2025, August 12). Socio-Legal Studies Association. https://www.slsa.ac.uk/post/can-maritime-law-inform-space-law-in-addressing-international-commons

Chalapathy, R., & Chawla, S. (2019). *Deep Learning for Anomaly Detection: A Survey* (arXiv:1901.03407). arXiv. https://doi.org/10.48550/arXiv.1901.03407

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

Chernavskikh, V., & Palayer, J. (n.d.). *Impact of Military Artificial Intelligence on Nuclear Escalation Risk*.

China and space: How space technologies boost China's intelligence capabilities as part of hybrid threats. (n.d.). *Hybrid CoE - The European Centre of Excellence for Countering Hybrid Threats*. Retrieved 20 May 2026, from https://www.hybridcoe.fi/publications/hybrid-coe-paper-21-china-and-space-how-space-technologies-boost-chinas-intelligence-capabilities-as-part-of-hybrid-threats/

China Practicing Satellite Warfare in Space. (n.d.). *Militarnyi*. Retrieved 25 May 2026, from https://militarnyi.com/en/news/china-practicing-satellite-warfare-in-space/

Chokepoint Above and Below the Surface: The Red Sea's Emerging Infrastructure Challenge | Istituto Affari Internazionali. (2025, December 13). https://www.iai.it/en/publications/c41/chokepoint-above-and-below-surface-red-seas-emerging-infrastructure-challenge

Cloud Computing Solutions for the Marine Corps: An Architecture to Support Expeditionary Logistics. (n.d.).

Coito, J. (2025). *Protecting Subsea Cables: Detect to Deter, Sue to Secure*. https://www.csis.org/analysis/protecting-subsea-cables-detect-deter-sue-secure

Conversation, T. (2021, September 21). *Undersea internet cables are causing geopolitical tension*. TNW | Insider. https://thenextweb.com/news/undersea-internet-cables-geopolitical-tension-syndication

Countering hybrid threats. (n.d.). Site Name Seo. Retrieved 4 May 2026, from https://www.nato.int/en/what-we-do/deterrence-and-defence/countering-hybrid-threats

Countering Ubiquitous Technical Surveillance. (n.d.). CIS. Retrieved 14 May 2026, from https://www.cisecurity.org/insights/white-papers/

Cross-cutting technologies in Chinese space activities: Raising the risk of hybrid threats. (n.d.). *Hybrid CoE - The European Centre of Excellence for Countering Hybrid Threats*. Retrieved 20 May 2026, from https://www.hybridcoe.fi/publications/hybrid-coe-paper-22-cross-cutting-technologies-in-chinese-space-activities-raising-the-risk-of-hybrid-threats/

DARPA's explainable AI (XAI) program: A retrospective—Gunning—2021—Applied AI Letters—Wiley Online Library. (n.d.). Retrieved 20 May 2026, from https://onlinelibrary.wiley.com/doi/full/10.1002/ail2.61

Digital twins and the future of maritime navigation. (n.d.). ADMIRALTY. Retrieved 5 May 2026, from https://www.admiralty.co.uk/about-us/case-study/digital-twin

Digital twins: Past, present and future | Scientific Reports. (n.d.). Retrieved 12 May 2026, from https://www.nature.com/articles/s41598-026-45272-z

Dugger, A. T. (2025, February 14). *Space as a Gray Zone: The Future of Orbital Warfare - Modern War Institute*. https://mwi.westpoint.edu/space-as-a-gray-zone-the-future-of-orbital-warfare/

Earth Science Data Systems, N. (2021, May 18). *Synthetic Aperture Radar (SAR) | NASA Earthdata* [Data Basics]. Earth Science Data Systems, NASA. https://www.earthdata.nasa.gov/learn/earth-observation-data-basics/sar

Edler, J. (n.d.). *TECHNOLOGY SOVEREIGNTY OF THE EU: NEEDS, CONCEPTS, PITFALLS AND WAYS FORWARD*.

Endsley, M. (1995). Endsley, M.R.: Toward a Theory of Situation Awareness in Dynamic Systems. Human Factors Journal 37(1), 32-64. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *37*, 32–64. https://doi.org/10.1518/001872095779049543

Estimating Worst-Case Frontier Risks of Open-Weight LLMs. (n.d.). Retrieved 19 June 2026, from https://arxiv.org/html/2508.03153v1

EU Maritime Security Strategy | EEAS. (n.d.). Retrieved 4 May 2026, from https://www.eeas.europa.eu/eeas/eu-maritime-security-strategy_en

Felleman, P. (1967, February 6). Hybrid simulation of the Apollo guidance navigation and control system. *Simulation and Support Conference*. Simulation and Support Conference. https://doi.org/10.2514/6.1967-233

Fernando, T., Denman, S., Sridharan, S., & Fookes, C. (2018). *GD-GAN: Generative Adversarial Networks for Trajectory Prediction and Group Detection in Crowds* (arXiv:1812.07667). arXiv. https://doi.org/10.48550/arXiv.1812.07667

Ferrari, S., & Cribari-Neto, F. (2004). Beta Regression for Modelling Rates and Proportions. *Journal of Applied Statistics*, *31*(7), 799–815. https://doi.org/10.1080/0266476042000214501

Flowers, A. (2025, May 8). Grey Zone Maritime Threats in European Waters & Strategies. *3GIMBALS*. https://3gimbals.com/insights/grey-zone-maritime-threats-in-european-waters-demand-new-security-strategies/

Fournier, M., Casey Hilliard, R., Rezaee, S., & Pelot, R. (2018). Past, present, and future of the satellite-based automatic identification system: Areas of applications (2004–2016). *WMU Journal of Maritime Affairs*, *17*(3), 311–345. https://doi.org/10.1007/s13437-018-0151-6

Full_journal_web_06-11.pdf. (n.d.). Retrieved 20 May 2026, from https://stratcomcoe.org/pdfjs/?file=/publications/download/full_journal_web_06-11.pdf?zoom=page-fit

G1115 Preparing for an IMO Member State Audit Scheme (IMSAS) on VTS. (n.d.). *IALA*. Retrieved 4 May 2026, from https://www.iala.int/product/g1115/

GAMBIT 1 KH-7 Film Recovery Vehicle. (n.d.). National Museum of the United States Air Force. Retrieved 29 April 2026, from https://www.nationalmuseum.af.mil/Visit/Museum-Exhibits/Fact-Sheets/Display/Article/195925/gambit-1-kh-7-film-recovery-vehicle/

GPS jamming of leader's plane puts Putin's hybrid warfare on Europe's radar. (2025, September 2). NBC News. https://www.nbcnews.com/world/europe/gps-jamming-von-der-leyen-plane-russia-hybrid-warfare-europe-ukraine-rcna228479

GPS Performances—Navipedia. (n.d.). Retrieved 11 April 2026, from https://gssc.esa.int/navipedia/index.php/GPS_Performances

Gunning, D., Vorm, E. s, Wang, J., & Turek, M. (2021). DARPA 's explainable AI (XAI) program: A retrospective. *Applied AI Letters*, *2*. https://doi.org/10.1002/ail2.61

Hamilton, W. L., Ying, R., & Leskovec, J. (2018). *Inductive Representation Learning on Large Graphs* (arXiv:1706.02216). arXiv. https://doi.org/10.48550/arXiv.1706.02216

Hanson, R. (2007). The Policy Analysis Market (A Thwarted Experiment in the Use of Prediction Markets for Public Policy). *Innovations: Technology, Governance, Globalization*, *2*(3), 73–88. https://doi.org/10.1162/itgg.2007.2.3.73

Håvard Tveit Ihle. (2026, May 28). *How far behind are open models?* [Online post]. LessWrong. https://www.lesswrong.com/posts/rJcCrXyEsJKmmDpWG/how-far-behind-are-open-models

Hayes, M. A., & Capretz, M. A. (2015). Contextual anomaly detection framework for big sensor data. *Journal of Big Data*, *2*(1), 2. https://doi.org/10.1186/s40537-014-0011-y

Helmus, T. C., Romita Grocholski, K., Liggett, T., Rhoades, A. L., Savitz, S., & Palmer, K. (2024). *Understanding and Countering China's Maritime Gray Zone Operations*. https://www.rand.org/pubs/research_reports/RRA2954-1.html

Hern, A. (2018, January 28). Fitness tracking app Strava gives away location of secret US army bases. *The Guardian*. https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases

Heuer, R. J. (with Center for the Study of Intelligence). (1999). *Psychology of intelligence analysis*. Center for the Study of Intelligence, Central Intelligence Agency.

Hicks, K. H. (with Friend, A. H.). (2019). *By Other Means Part I: Campaigning in the Gray Zone* (1st ed). Center for Strategic & International Studies.

Hitchens, T. (2025, March 18). 5 Chinese satellites practiced 'dogfighting' in space, Space Force says. *Breaking Defense*. https://breakingdefense.com/2025/03/5-chinese-satellites-practiced-dogfighting-in-space-space-force-says/

Horlings, T., Lindelauf, R., & Rietjens, S. (2023). *CHAPTER 10 Battling information overload in military intelligence & security organisations* (pp. 221–236). https://doi.org/10.24415/9789400604537-014

Huawei bans leave fiber operators little choice besides Nokia. (n.d.). Light Reading. Retrieved 21 May 2026, from https://www.lightreading.com/fttx/huawei-bans-leave-fiber-operators-with-little-choice-besides-nokia

Hwang, H., & Joe, I. (2023). Enhancing IoT Connectivity and Services for Worldwide Ships through Multi-Region Fog Cloud Architecture Platforms. *Electronics*, *12*(20), 4250. https://doi.org/10.3390/electronics12204250

Hybrid Threats, A Strategic Communications Perspective. (n.d.). NATO Strategic Command. Retrieved 20 May 2026, from https://stratcomcoe.org/pdfjs/?file=/publications/download/2nd_book_short_digi_pdf.pdf?zoom=page-fit

Iii, J. M. M., & Jensen, B. (2025). *Signals in the Swarm: The Data Behind China's Maritime Gray Zone Campaign Near Taiwan*. https://www.csis.org/analysis/signals-swarm-data-behind-chinas-maritime-gray-zone-campaign-near-taiwan

Increasing maritime activity in the Strait of Gibraltar. (n.d.). *HZ CONTAINERS.Com*. Retrieved 17 January 2026, from https://hz-containers.com/en/news/increasing-maritime-activity-in-the-strait-of-gibraltar/

January 20, D. M. • & 2026. (n.d.). *Europe Readies Law to Eject Chinese Equipment From Telecoms*. Retrieved 21 May 2026, from https://www.bankinfosecurity.com/europe-readies-law-to-eject-chinese-equipment-from-telecoms-a-30566

Jensen, B. (2025). *How to Exorcise Russia's Ghost Fleet*. https://www.csis.org/analysis/how-exorcise-russias-ghost-fleet

Jiang, B., Dong, X., Deng, M., Wan, F., Wang, T., Li, X., Zhang, G., Cheng, Q., & Lv, S. (2023). Geolocation Accuracy Validation of High-Resolution SAR Satellite Images Based on the Xianning Validation Field. *Remote Sensing*, *15*(7), 1794. https://doi.org/10.3390/rs15071794

Jr, F. E. W. (n.d.). *Data Fusion Group Perspective*.

Karamolegkos, S., Dourvas, N., Episkopos, N., Pikounis, K., Michail, E., Ioannidis, K., & Vrochidis, S. (2025). SIAP: Synthetic dataset for maritime vessel risk profiling and illegal activity prediction. *Data in Brief*, *63*, 112101. https://doi.org/10.1016/j.dib.2025.112101

Khodabakhshian, A., Re Cecconi, F., & Droguett, E. (2025). Probabilistic risk identification and assessment model for construction projects using elicitation based bayesian network. *Journal of Information Technology in Construction*, *30*, 185–212. https://doi.org/10.36680/j.itcon.2025.009

Komiss, W., & Huntzinger, L. (n.d.). *The Economic Implications of Disruptions to Maritime Oil Chokepoints*.

Laxhammar, R. (2008). *Anomaly detection for sea surveillance*. 1–8.

Li, G., Zhang, X., Shu, Y., Wang, C., Guo, W., & Wang, J. (2024). Ship Anomalous Behavior Detection in Port Waterways Based on Text Similarity and Kernel Density Estimation. *Journal of Marine Science and Engineering*, *12*(6), 968. https://doi.org/10.3390/jmse12060968

Looney, R. E. (2004). DARPA's Policy Analysis Market for Intelligence: Outside the Box or Off the Wall? *International Journal of Intelligence and CounterIntelligence*, *17*(3), 405–419. https://doi.org/10.1080/08850600490446745

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2019). *Explainable AI for Trees: From Local Explanations to Global Understanding* (arXiv:1905.04610). arXiv. https://doi.org/10.48550/arXiv.1905.04610

Lytra, I., Vidal, M.-E., Orlandi, F., & Attard, J. (2017, June 21). *A Big Data Architecture for Managing Oceans of Data and Maritime Applications*. Zenodo. International Conference on Engineering, Technology and Innovation (ICE). https://doi.org/10.5281/zenodo.815361

MacColl, N. B., Pia Huesch, Jamie. (2025, July 21). European critical infrastructure still struggles with Chinese ICT vendors. *The Strategist*. https://www.aspistrategist.org.au/european-critical-infrastructure-still-struggles-with-chinese-ict-vendors/

Maintaining the UK's Intelligence Edge in the Grey Zone. (2026, January 16). https://www.rusi.orghttps://www.rusi.org

Maliarchuk, T., Danyk, Y., & Briggs, C. (2019). Hybrid Warfare and Cyber Effects in Energy Infrastructure. *Connections: The Quarterly Journal*, *18*(1). https://connections-qj.org/article/hybrid-warfare-and-cyber-effects-energy-infrastructure

Mallory, K. (2018). *New Challenges in Cross-Domain Deterrence*. https://www.rand.org/pubs/perspectives/PE259.html

Maneuvering target tracking. (n.d.). Retrieved 11 April 2026, from https://www.academia.edu/74746433/Maneuvering_target_tracking

Maritime Anomaly Detection using Density-based Clustering and Recurrent Neural Network | The Journal of Navigation | Cambridge Core. (n.d.-a). Retrieved 29 April 2026, from https://www.cambridge.org/core/journals/journal-of-navigation/article/abs/maritime-anomaly-detection-using-densitybased-clustering-and-recurrent-neural-network/974C7AEF33E437D0059D517183610709

Maritime Anomaly Detection using Density-based Clustering and Recurrent Neural Network | The Journal of Navigation | Cambridge Core. (n.d.-b). Retrieved 29 April 2026, from https://www.cambridge.org/core/journals/journal-of-navigation/article/abs/maritime-anomaly-detection-using-densitybased-clustering-and-recurrent-neural-network/974C7AEF33E437D0059D517183610709

Maritime—Deimos Space. (2023, May 29). https://deimos-space.com/maritime-2/

McKinley, C. A. (n.d.). *The Guardians of Space: Organizing America's Space Assets for the Twenty-First Century*.

Mediterranean ports, hubs of global trade. (2023, May 3). *We Build Value*. https://www.webuildvalue.com/en/reportage/mediterranean-ports-global-trade.html

MIĘTKIEWICZ, R. (2025). *Hybrid threats in the Baltic Sea. The results of analysis of countermeasure options*. Polish Naval Academy of the Heroes of Westerplatte.

Milledge, T. J., Pradhan, B., & Shukla, N. (2026). *Techno-economic modelling of fleet architectures for undersea surveillance*. https://opus.lib.uts.edu.au/handle/10453/193483

Monaghan, S., Svendsen, O., Darrah, M., & Arnold, E. (2023). *NATO's Role in Protecting Critical Undersea Infrastructure*. https://www.csis.org/analysis/natos-role-protecting-critical-undersea-infrastructure

Moratelli Junior, L. (2014). *Integrated controller based on hybrid concept and geometric control for dynamically positioned vessels*. https://doi.org/10.11606/T.3.2016.tde-13072016-143252

NATO. (n.d.). *Defence expenditures and NATO's 5% commitment*. NATO. Retrieved 13 November 2025, from https://www.nato.int/cps/en/natohq/topics_49198.htm

NATO Review—Fortifying the Baltic Sea—NATO's defence and deterrence strategy for hybrid threats. (2025, May 5). NATO Review. https://www.nato.int/docu/review/articles/2025/05/05/fortifying-the-baltic-sea-natos-defence-and-deterrence-strategy-for-hybrid-threats/index.html

Navies are Our Guide for Space Strategy—Center for Maritime Strategy. (2025, July 22). https://centerformaritimestrategy.org/publications/navies-are-our-guide-for-space-strategy/

NSO NSDD. (n.d.). Retrieved 10 June 2026, from https://nso.nato.int/nso/nsdd/main/standards/ap-details/2272/EN

Ören, A. (2026). The future of space Security: Toward NATO's evolving functional areas in the era of hybrid warfare. *Space Policy*, 101746. https://doi.org/10.1016/j.spacepol.2026.101746

P, S. R. (2024, June 22). Types of Satellites in a glimpse. *Hex-Star Universe*. https://hexstaruniverse.com/comparision-between-types-of-satellites/

Palacios, A. (2025, May 28). *La Armada activa una misión de inteligencia ante el paso de buques de guerra rusos por el estrecho de Gibraltar*. El Debate. https://www.eldebate.com/espana/defensa/20250528/armada-activa-mision-inteligencia-ante-paso-buques-guerra-rusos-estrecho-gibraltar_301679.html

Palak, R., & Nguyen, N. T. (2017). *Prediction markets as a vital part of collective intelligence*. 137–142. https://doi.org/10.1109/INISTA.2017.8001146

Pallotta, G., Vespe, M., & Bryan, K. (2013). Vessel Pattern Knowledge Discovery from AIS Data: A Framework for Anomaly Detection and Route Prediction. *Entropy*, *15*, 2218–2245. https://doi.org/10.3390/e15062218

(PDF) A Deep Learning Model for Ship Trajectory Prediction Using Automatic Identification System (AIS) Data. (n.d.). Retrieved 4 May 2026, from https://www.researchgate.net/publication/369712377_A_Deep_Learning_Model_for_Ship_Trajectory_Prediction_Using_Automatic_Identification_System_AIS_Data

(PDF) A Multiscale Anomaly Detection Framework for AIS Trajectories via Heat Graph Laplacian Diffusion. (n.d.). *ResearchGate*. https://doi.org/10.23919/EUSIPCO63174.2024.10715305

(PDF) A Multi-Task Deep Learning Architecture for Maritime Surveillance Using AIS Data Streams. (n.d.). Retrieved 14 May 2026, from https://www.researchgate.net/publication/330878370_A_Multi-Task_Deep_Learning_Architecture_for_Maritime_Surveillance_Using_AIS_Data_Streams

(PDF) An Automatic Identification System (AIS) Database for Maritime Trajectory Prediction and Data Mining. (n.d.). ResearchGate. https://doi.org/10.48550/arXiv.1607.03306

(PDF) Anomaly Detection in Maritime Domain based on Spatio-Temporal Analysis of AIS Data Using Graph Neural Networks. (n.d.). *ResearchGate*. https://doi.org/10.1109/ICVISP54630.2021.00033

(PDF) Attention-Aware Graph Neural Network Modeling for AIS Reception Area Prediction. (n.d.). Retrieved 4 May 2026, from https://www.researchgate.net/publication/396437311_Attention-Aware_Graph_Neural_Network_Modeling_for_AIS_Reception_Area_Prediction

(PDF) Data Fusion for Wide-Area Maritime Surveillance. (n.d.). *ResearchGate*. Retrieved 4 May 2026, from https://www.researchgate.net/publication/259772622_Data_Fusion_for_Wide-Area_Maritime_Surveillance

(PDF) EO-ALERT: Next Generation Satellite Processing Chain for Security-driven Early Warning Capacity in Maritime Surveillance and Extreme Weather Events. (n.d.). ResearchGate. https://doi.org/10.5281/zenodo.3243836

(PDF) Extracting Maritime Traffic Networks from AIS Data Using Evolutionary Algorithm. (n.d.). *ResearchGate*. https://doi.org/10.1007/s12599-020-00661-0

(PDF) GeoTrackNet-A Maritime Anomaly Detector using Probabilistic Neural Network Representation of AIS Tracks and A Contrario Detection. (n.d.). ResearchGate. https://doi.org/10.48550/arXiv.1912.00682

(PDF) Graph-based anomaly detection. (n.d.). *ResearchGate*. https://doi.org/10.1145/956750.956831

(PDF) Machine Learning-Assisted Anomaly Detection in Maritime Navigation Using AIS Data. (n.d.). *ResearchGate*. https://doi.org/10.1109/PLANS46316.2020.9109806

(PDF) Maritime Clouds for the European Ports. (n.d.). *ResearchGate*. https://doi.org/10.1109/PCi.2012.39

(PDF) Spatio-Temporal Graphs Beyond Grids: Benchmark for Maritime Anomaly Detection. (2025, December 28). ResearchGate. https://doi.org/10.48550/arXiv.2512.20086

Pérez Carrasco, F. J., Fernández García, A., García García, A., Ruiz Bejerano, V., Gutiérrez Martín, Á., & Belmonte Hernández, A. (2025). Cloud-based architecture for hydrophone data acquisition and processing of surface and underwater vehicle detection. *Journal of Marine Science and Engineering*, *13*(8), 1455.

Philippines accuses China's forces of harassing fisheries vessels in the South China Sea. (2025, January 25). AP News. https://apnews.com/article/south-china-sea-sandy-cay-philippines-606e460688e43f5fecc5264ab315feb9

Popular Orbits 101. (2017, November 30). *Aerospace Security*. https://aerospace.csis.org/aerospace101/earth-orbit-101/

Pražák, J. (2022). On the Threshold of Space Warfare. *Astropolitics*, *20*(2–3), 175–191. https://doi.org/10.1080/14777622.2022.2142351

Press Release: Indo-Pacific Commander Calls Bitcoin A Tool For U.S. 'Power Projection' In Senate Testimony | Bitcoin Policy Institute. (n.d.). Retrieved 23 May 2026, from https://www.btcpolicy.org/articles/press-release-indo-pacific-commander-calls-bitcoin-a-tool-for-us-power-projection-in-senate-testimony

published, M. T. (2025, August 13). *Crew of Russian 'shadow fleet' tanker charged with intentionally sabotaging undersea cables—Finnish authorities claim ship dragged its anchor 56 miles along the sea floor, causing $70 million in damage to five cables*. Tom's Hardware. https://www.tomshardware.com/networking/crew-of-russian-shadow-fleet-tanker-charged-with-intentionally-sabotaging-undersea-cables-finnish-authorities-claim-ship-dragged-its-anchor-56-miles-along-the-sea-floor-causing-usd70-million-in-damage-to-five-cables

Qin, B., Zhu, H., Zhu, R., Zhao, M., Qiu, M., & Li, Q. (2025). Space-to-ground infrared camouflage with radiative heat dissipation. *Light: Science & Applications*, *14*(1), 137. https://doi.org/10.1038/s41377-025-01824-y

Revised Maritime Spacepower Theory. (2023, July 31). Air University (AU). https://www.airuniversity.af.edu/JIPA/Display/Article/3475007/revised-maritime-spacepower-theory/

Riveiro, M., Pallotta, G., & Vespe, M. (2018). Maritime anomaly detection: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *8*, e1266. https://doi.org/10.1002/widm.1266

Ross, S. M. (2014). *Introduction to probability models* (Eleventh edition). Elsevier.

Running Into Open Secrets: How to Investigate Using the Strava Fitness App – Global Investigative Journalism Network. (n.d.). Retrieved 14 May 2026, from https://gijn.org/stories/investigations-using-strava-fitness-app/

Rusia incrementa sus movimientos navales por el estrecho de Gibraltar. (n.d.). Maribero. Retrieved 17 January 2026, from https://www.maribero.com/article/zn87hee4t8c3yikntiw8rkf0

Santamaria, C., Alvarez, M., Greidanus, H., Syrris, V., Soille, P., & Argentieri, P. (2017). Mass Processing of Sentinel-1 Images for Maritime Surveillance. *Remote Sensing*, *9*, 678. https://doi.org/10.3390/rs9070678

Secure World Foundation: 2026 Global Counterspace Capabilities Report. (n.d.). Retrieved 25 May 2026, from https://www.swfound.org/publications-and-reports/2026-global-counterspace-capabilities-report

Shabbir, M. O., & Omer, R. (2026). Leading Security Risks Originating in Space Hybrid Operations. *Research Journal for Social Affairs*, *4*(3(s2)), 161–166. https://doi.org/10.71317/RJSA.004.03(s2).0840

Sheppard, L. R., & Conklin, M. (2019). *Warning for the Gray Zone*. https://www.csis.org/analysis/warning-gray-zone

Soldi, G., Gaglione, D., Forti, N., Simone, A. D., Daffinà, F. C., Bottini, G., Quattrociocchi, D., Millefiori, L. M., Braca, P., Carniel, S., Willett, P., Iodice, A., Riccio, D., & Farina, A. (2021a). Space-based Global Maritime Surveillance. Part I: Satellite Technologies. *IEEE Aerospace and Electronic Systems Magazine*, *36*(9), 8–28. https://doi.org/10.1109/MAES.2021.3070862

Soldi, G., Gaglione, D., Forti, N., Simone, A. D., Daffinà, F. C., Bottini, G., Quattrociocchi, D., Millefiori, L. M., Braca, P., Carniel, S., Willett, P., Iodice, A., Riccio, D., & Farina, A. (2021b). Space-based Global Maritime Surveillance. Part II: Artificial Intelligence and Data Fusion Techniques. *IEEE Aerospace and Electronic Systems Magazine*, *36*(9), 30–42. https://doi.org/10.1109/MAES.2021.3070884

SOTDMA. (n.d.). Retrieved 30 March 2026, from http://www.allaboutais.com/index.php/en/technical-info/transmission-types/105-ais-technicals/technical-lfundamentals/93-sotdma

Stanciuc, A.-M. (2026, May 4). *Brussels reissues its Huawei warning, and prepares to make it stick*. TNW | Eu. https://thenextweb.com/news/eu-huawei-zte-connectivity-infrastructure-recommendation

Starboard Maritime Intelligence secures NZ$23M Series A to scale global AI-powered maritime intelligence technology | Starboard Maritime Intelligence. (n.d.). Retrieved 10 March 2026, from https://www.starboardintelligence.org/press/starboard-maritime-intelligence-secures-nz-23m-series-a-to-scale-global-ai-powered-maritime-intelligence-technology.html

Sun, H., & Qin, Y. (2022). Stealthy Configuration Optimization Design and RCS Characteristics Study of Microsatellite. *Aerospace*, *9*(12), 815. https://doi.org/10.3390/aerospace9120815

Swords and Shields: Navigating the Modern Intelligence Landscape. (n.d.). ORF Middle East. Retrieved 14 May 2026, from https://orfme.org/research/swords-and-shields-navigating-the-modern-intelligence-landscape/

Tactical And Strategic Missile Guidance (2012). (n.d.). Retrieved 11 April 2026, from http://archive.org/details/TacticalAndStrategicMissileGuidance2012

Taipei, K. H. in, York, S. L. in N., Joiner, S., Arenas, I. de la T., & London, D. C. in. (2026a, April 27). *Inside China's plans to fight in space*. https://ig.ft.com/space-weapons

Taipei, K. H. in, York, S. L. in N., Joiner, S., Arenas, I. de la T., & London, D. C. in. (2026b, April 27). *Inside China's plans to fight in space*. https://ig.ft.com/space-weapons

Thadani, A., & Allen, G. C. (2023). *Mapping the Semiconductor Supply Chain: The Critical Role of the Indo-Pacific Region*. https://www.csis.org/analysis/mapping-semiconductor-supply-chain-critical-role-indo-pacific-region

The Changing Character Of Maritime Choke Points: Who Throttles Whom? (n.d.). Hoover Institution. Retrieved 20 May 2026, from https://www.hoover.org/research/changing-character-maritime-choke-points-who-throttles-whom

The geopolitics of AI and the rise of digital sovereignty. (n.d.). *Brookings*. Retrieved 20 May 2026, from https://www.brookings.edu/articles/the-geopolitics-of-ai-and-the-rise-of-digital-sovereignty/

The Shadow Fleet: How Russian Oil Flows Through Georgian Ports. (2026, January 27). iFact. https://ifact.ge/en/the-shadow-fleet-how-russian-oil-flows-through-georgian-ports/

The U.S. Coast Guard and the U.S. Space Force: An Essential Partnership in Maritime Security. (n.d.). *USNI Blog*. Retrieved 25 May 2026, from https://blog.usni.org/posts/2020/10/13/the-u-s-coast-guard-and-the-u-s-space-force-an-essential-partnership-in-maritime-security

'They will be back': How China's 'dark' fleets are plundering the world's oceans. (2020, December 18). *ABC News*. https://www.abc.net.au/news/2020-12-19/how-china-is-plundering-the-worlds-oceans/12971422

Tokmetzis, D., & Schohaus, B. (2024, June 20). *Almost 200 Russian ships suspected of spying in the North Sea*. Follow the Money - Platform for Investigative Journalism. https://www.ftm.eu/articles/north-sea-1-suspicious-russian-ships-in-european-waters-and-havens

Too Much Information: Ineffective Intelligence Collection. (2019, August 18). Harvard International Review. https://hir.harvard.edu/too-much-information/

Tran, M., Shipard, J., Mulyono, H., Wiliem, A., & Fookes, C. (2023). *SafeSea: Synthetic Data Generation for Adverse & Low Probability Maritime Conditions* (arXiv:2311.14764). arXiv. https://doi.org/10.48550/arXiv.2311.14764

Types of Earth Orbits and Their Applications—eoPortal. (n.d.). Retrieved 9 April 2026, from https://www.eoportal.org/other-space-activities/orbit-types

Ulrichsen, K. C., & Krane, J. (n.d.). *Maritime Chokepoints and Risks to Global Shipping and Energy Security*.

Un momento…. (n.d.-a). Retrieved 12 February 2026, from https://www.researchgate.net/profile/Kjell-Hausken/publication/309088927_Cost_benefit_analysis_of_war/links/59d9ebff458515a5bc2b1ce7/Cost-benefit-analysis-of-war.pdf?_tp=eyJjb250ZXh0Ijp7ImZpcnN0UGFnZSI6InB1YmxpY2F0aW9uIiwicGFnZSI6InB1YmxpY2F0aW9uIn19&__cf_chl_tk=1tWaHG00LJQgsQx8obQ7TWStu2tSgnmFtjStPW4emEQ-1770936255-1.0.1.1-DLus2RIqGrFmT6LE0vTJ0WIcL0GWfFZnYWLyO3peLqM

Un momento…. (n.d.-b). Retrieved 29 April 2026, from https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/widm.1266

UW ADINT: Advertising as Surveillance. (n.d.). Retrieved 14 May 2026, from https://adint.cs.washington.edu/

Velasquez, V. G. J., Michael Kidd, Giana Farry, Mohamed Abdelsalam, Aakash Palathra, Andrew. (2025, December 1). Strategic Relevance Of Maritime Chokepoints For Global Trade. *The Defence Horizon Journal*. https://tdhj.org/blog/post/nato-maritime-chokepoints/

Venskus, J., Treigys, P., Bernatavičienė, J., Tamulevičius, G., & Medvedev, V. (2019). Real-Time Maritime Traffic Anomaly Detection Based on Sensors and History Data Embedding. *Sensors*, *19*(17), 3782. https://doi.org/10.3390/s19173782

Vincent, A. (2015, September 18). *Exhausting Big Data to Defend against Cyber Security Threats: A counter-intelligence approach.*

Vines, P., Roesner, F., & Kohno, T. (2017). Exploring ADINT: Using Ad Targeting for Surveillance on a Budget - or - How Alice Can Buy Ads to Track Bob. *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society*, 153–164. https://doi.org/10.1145/3139550.3139567

Vogt, J. (2026, June 9). *Data Center Warfare: Defending the Key Terrain of AI Infrastructure - Modern War Institute*. https://mwi.westpoint.edu/data-center-warfare-defending-the-key-terrain-of-ai-infrastructure/

Wall, C., & Morcos, P. (2021). *Invisible and Vital: Undersea Cables and Transatlantic Security*. https://www.csis.org/analysis/invisible-and-vital-undersea-cables-and-transatlantic-security

Wang, X., & Xiao, Y. (2023). A Deep Learning Model for Ship Trajectory Prediction Using Automatic Identification System (AIS) Data. *Information*, *14*, 212. https://doi.org/10.3390/info14040212

Weigle, B. (2007). *Prediction Markets: Another Tool in the Intelligence Kitbag*.

What is Remote Sensing? (n.d.). *Windward*. Retrieved 13 April 2026, from https://windward.ai/glossary/what-is-remote-sensing/

Why U.S. and Chinese satellites are 'dogfighting' in orbit. (2025, December 18). The Washington Post. https://www.washingtonpost.com/technology/interactive/2025/space-military-satellite-china-united-states/

Why we have six months to regulate the AI that Europe isn't building. (2026, June 17). *CEPS*. https://www.ceps.eu/why-we-have-six-months-to-regulate-the-ai-that-europe-isnt-building/

Wiesner, P., O'Neill, D. W., Larosa, F., & Kao, O. (2025). *Efficiency Will Not Lead to Sustainable Reasoning AI* (arXiv:2511.15259). arXiv. https://doi.org/10.48550/arXiv.2511.15259

Winkel, J., Willems, T., O'Driscoll, C., & Fernandez-Hernandez, I. (2026). *SeaSpoofFinder—Potential GNSS Spoofing Event Detection Using AIS* (arXiv:2602.16257; Version 1). arXiv. https://doi.org/10.48550/arXiv.2602.16257

World Meteorological Organization (Ed.). (1998). *Guide to wave analysis and forecasting* (2nd ed). Secretariat of the World Meteorological Organization.

Wright, N. (n.d.). *Human, Machine, War: How the Mind-Tech Nexus will Win Future Wars*.

XAI | DARPA. (n.d.). Retrieved 20 May 2026, from https://www.darpa.mil/research/programs/explainable-artificial-intelligence

Yeh, P. F. (2006). *Using Prediction Markets to Enhance US Intelligence Capabilities: (741212011-004)* [Data set]. https://doi.org/10.1037/e741212011-004

Yelland, M. J., Moat, B. I., Taylor, P. K., Pascal, R. W., Hutchings, J., & Cornell, V. C. (1998). Wind Stress Measurements from the Open Ocean Corrected for Airflow Distortion by the Ship. *Journal of Physical Oceanography*, *28*(7), 1511–1526. https://doi.org/10.1175/1520-0485(1998)028%253C1511:WSMFTO%253E2.0.CO;2

Zhang, B., Hirayama, K., Ren, H., Wang, D., & Li, H. (2023). Ship Anomalous Behavior Detection Using Clustering and Deep Recurrent Neural Network. *Journal of Marine Science and Engineering*, *11*(4), 763. https://doi.org/10.3390/jmse11040763

Zhou, F., Yu, K., Xie, W., Lyu, J., Zheng, Z., & Zhou, S. (2024). Digital Twin-Enabled Smart Maritime Logistics Management in the Context of Industry 5.0. *IEEE Access*, *12*, 10920–10931. https://doi.org/10.1109/ACCESS.2024.3354838

(N.d.).