AI-ARC Baltic Demo: Detecting Illegal Activities at Sea

Pontus Svenson*, Anders Holst*, Anders Wallberg*, Paavo Nevalainen†, Farshad Farahnakian†, Alfonso Álamo‡, Vincenzo Germinara‡, Daniel Schweizer§, Matthis Leicht¶, Mathias Anneken¶, Adrian H. Hoppe¶, Aristeidis Karalis‖ Ashraf Labib‖, Maria Eugenia Beltrán**, Liss Hernández**, Petteri Partanen††, Minna Markkanen††

*RISE Research Institutes of Sweden firstname.lastname@ri.se
† University of Turku firstname.lastname@utu.fi
‡ TREE Technology firstname.lastname@treelogic.com
§ Fraunhofer EMI Daniel.Schweizer@emi.fraunhofer.de
¶ Fraunhofer IOSB firstname.lastname@iosb.fraunhofer.de
‖University of Portsmouth firstname.lastname@port.ac.uk
** Universidad Politécnica de Madrid {mebeltran,lhernandez}@lst.tfo.upm.es
††Laurea University of Applied Science firstname.lastname@laurea.fi

***Abstract*—We describe the AI-ARC (Artificial Intelligence-based Virtual Control Room for the Arctic) system, which aims to enhance maritime domain awareness and surveillance. The system is micro-service based and fuses data from various sources, utilizing AI-driven micro-services and an advanced visualization platform to increase the situation awareness of maritime surveillance operators. The results of the Baltic sea demonstration, aiding in the detection of illegal activities, environmental protection, are presented. The system was evaluated using historical data from real criminal incidents. The results show that the AI-ARC approach could help increase the situation awareness of law enforcement operators.**

***Index Terms*—situation awareness, smuggling, infrastructure protection, maritime domain awareness, anomaly detection, intent detection**

I. Introduction

Maritime domain awareness (MDA) has long been an important research area [1]. Due to the natural growth of global traffic and also because of recent geopolitical shifts, the importance of maritime surveillance has increased manifold. Monitoring maritime activities helps identify suspicious vessels and potential security threats such as piracy, smuggling, and illegal fishing. By tracking ship movements, authorities can detect and respond to security risks in a timely manner, ensuring the safety of both maritime assets and coastal communities. Maritime surveillance also contributes to search, rescue and assist operations by providing real-time information on the location of vessels in distress. SAR (Synthetic Aperture Radar) images and AIS (Automatic Identification System) data enable authorities to coordinate rescue efforts more effectively, reducing response times and increasing the chances of saving lives in emergency situations.

Environmental protection and protection of scarce natural resources is also an important application area. By tracking fishing vessels and monitoring fishing activities, authorities can prevent over-fishing, protect endangered species, and promote the conservation of marine habitats. Detecting e.g. drifting ships or ships being anchored to dangerous position helps to prevent environmental hazards such as oil spills, illegal dumping, and pollution. By identifying polluters and enforcing environmental regulations, authorities can mitigate the impact of human activities on marine ecosystems and preserve biodiversity.

In this paper, we describe some of the results of the AI-ARC Horizon 2020 project. The project focused on developing a micro-service based system of systems that can help maritime surveillance operators increase their situation awareness. To do this, we developed and demonstrated a system of systems for acquiring data, fusing it to produce situation pictures and threat analyses, and then visualize it in 2D and 3D for the operators. The goals were to generate meaningful insights and predictions, and create actionable intelligence by using multiple data sources processed by combinations of AI-based micro-services from multiple providers. All visualized across different platforms in 2D & 3D using AI-ARC's Virtual Control Room.

To accomplish MDA, the research was organized around user needs as defined by stakeholder partners. An architecture that enables users to define different *service chains* for different use cases was implemented. Aggregating intention detection services then produce alerts warning the user about possible illegal activities.

Situational awareness is defined, following [2], as "the perception of environmental elements and events with respect to time or space, the comprehension of their meaning, and the projection of their future status". Situation awareness thus encompasses both the detection of relevant elements, the understanding of their meaning, and the (short-term) prediction into the future of how the situation will evolve. These three components of situation awareness are sometimes described as level 1 (perception), level 2 (comprehension), and level 3 (projection). Note that the sometimes-used term situational understanding [3] is the same as level 2 situation awareness. The AI-ARC services aim to contribute to all three levels of

This research was funded by the European Commission under grant no. 101021271 (Artificial Intelligence based Virtual Control Room for the Arctic (AI-ARC)).

situation awareness.

One effect of the improved situation awareness should be that the human operators can spend more time analysing and understanding the maritime situation, without having to look at all data and do routine analysis that the computer can do [4], [5]. By using the AI-ARC services to detect important events, the users get more time for the qualitative reasoning and analysis.

The main services developed in AI-ARC, covering levels 0, 1, 2, and 3 of the JDL model [6] are:

- Risk and reliability analysis
  - Risk Index Computation Service (developed by University of Portsmouth, UoP)
  - Reliability Assessment of ML-Services (developed by the Fraunhofer Institute for High-Speed Dynamics, FhG-EMI)
- Anomaly detection and intent recognition
  - DBN-Based Anomaly Detection (developed by the Faunhofer Institute of Optronics, System Technology, and Image Exploitation, FhG-IOSB)
  - Anomaly Detection and Intention Detection (developed by the Swedish Research Institutes, RISE)
  - Anomaly Detection on Earth Observation Data (developed by Thales Alenia Space, TAS)
  - Anomaly Detection in Vessel's Behaviour (developed by Tree Technology SA, TREE)
- Prediction of ice and vessel traffic
  - Satellite-Based Vessel / Iceberg Detection Service (developed by Telespzaio, TPZF)
  - Satellite-Based Sea Ice Coverage Maps Service (developed by Telespzaio, TPZF)
  - Detection of Ice Blocks (developed by Tree Technology SA, TREE)
  - Prediction of Icepacks (developed by University of Turku, UTU)
  - Vessel Traffic Prediction Service (developed by University of Turku, UTU)

In addition to the analysis services, user interaction (level 5 fusion [7]) was covered by four services:

- DigLT map display (deveoped by FhG-IOSB)
- In-Situ observation app (developed by FhG-IOSB)
- Big-data analytics dashboard (developed by Sampas, SAM)
- Business process modelling service (developed by SAM)

The DigLT is a system for distributed situation visualization and planning, both web-based and in virtual reality, in which any number of users can work independently or together on the same situation. It combines and displays all the results that are generated by the AI-ARC services.

AI-ARC performed two demonstrations and a pre-operational validation test period. In this paper, we focus on descring the AI-ARC Baltic demonstration, which focused on detecting illegal activities. Further details on the technical implementation of the services and the Arctic demonstration will be presented elsewhere.

This paper is outlined as follows. We first discuss the data sources and use cases chosen for the demo, followed by a description of the system architecture. We then turn to the services developed, after which the evaluation results are described.

II. Data sources and use cases

The main data source used in AI-ARC is AIS data. In addition, satellite images were used to detect unclear waters (a possible indication of environmental violations) and dark vessels, i.e., ships that do not have an active AIS transponder. For the development and Baltic demonstration, we choose to use real historical data as opposed to conducting live exercises with vessels moving at sea according to a script. The reason for this was to show performance with real incidents of criminal behaviour. In the AI-ARC Arctic demo, the capability of the AI-ARC system to handle live data was demonstrated, with a focus on safe navigation and search and rescue operations.

We now give a brief overview of the use cases.

Grounding: Show that it is possible to give warning when a vessel is about to ground. Using a two-layer approach: a risk index computation attempts to give an early alert; a grounding prediction that gives an alert when a vessel will ground in near-time. Shown using data from a real grounding taking place off the west coast of Sweden in September 2015

Environment: Demonstrate the capability of AI-ARC services to detect environmental crimes such as oil spills, illegal residue discharges, and to identify possible suspect vessels. This is achieved by: Satellite image processing to detect unclear waters; Satellite image processing to detect dark vessels; Detection of characteristic movements for tank cleaning; combining these weak signals in an aggregating detector.

Critical infrastructure protection: Is it possible to detect ships that move suspiciously near a critical underwater infrastructure? Approach: Detection of dark vessels from satellite image processing; Detection of vessels that loiter around the infrastructure (i.e., are stationary or moving only very little); Combining weak signals using a dynamic Bayesian network. To test this use case, we choose to investigate the area south of Sweden, east of Bornholm on 22nd September 2022, i.e., corresponding to the North Stream explosion.

Illegal fishing: Demonstrate the capability of detecting and reporting fishing in restricted areas or exclusive economic zones (EEZ). Approach: Detection of dark vessels (turned off AIS) from satellite image processing; Detection of vessels that show fishing movement patterns in predefined zones; Detection of vessel origin based on country code ; Combining weak signals using a dynamic Bayesian network.

Smuggling: Determine if it is possible to detect potential smuggling by analyzing AIS data. This is possibly the most difficult use case, since smuggling is a very heterogeneous task – it can happen in a large variety of ways. Given the high stakes, the perpetrators do everything to appear normal We have only access to AIS, so will not see any small, quick dark vessels involved. Approach: Focus on hand-over of goods at sea. This requires some potentially detectable activity

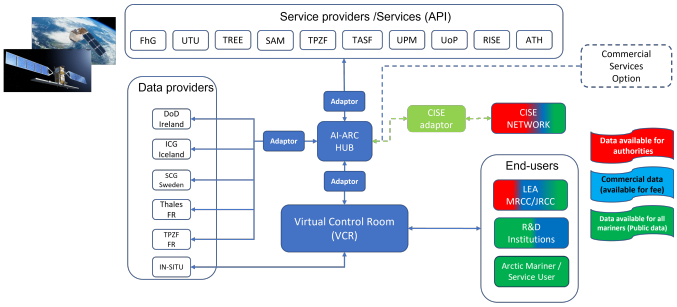

Fig. 1. Overview of the AI-ARC architecture

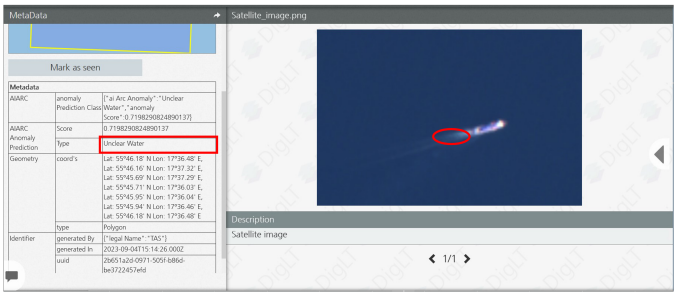

Fig. 2. Pollution detected by processing satellite image data, visualized in the DigLT. Metadata related to the detection can be shown.

during travel: Meetings between vessels at sea; Seemingly unmotivated route to destination (to get to a meeting place); Unusual maneauvres (while meeting someone, or picking something up). Smuggling analysis was demonstrated using real historical data from two cases: A suspected smuggling that occurred in January 2010 off the west coast of Sweden; A confirmed smuggling that occurred in February 2020 in Danish waters.

III. System Architecture

Figure 1 shows an overview of the system architecture.

The AI-ARC system of systems defines several kinds of messages that are exchanged between the services:

- An AnomalyPrediction should be displayed in the alarm list of the VCR and is serious enough that it warrants investigation by the operator.
- An AnomalyIndication message is not displayed in the alarm list and in general not investigated by the operator.
- An AnomalyIntention message is displayed in the alarm list and investigated by the operator as well as possibly by intelligence analysts. In the process of investigating these, the involved AnomalyPrediction and AnomalyIndication messages might also be explored by the user.

The AI-ARC AI services are based on series of basic micro-services developed during the project that are meant to be combined in different ways depending on the use cases. In general, the AI-ARC services consist of

- Data ingestion services that connect the AI-ARC ecosystem to data sources such as AIS track providers and satellite SAR / imagery. These are producing Vessel messages (formatted Tracks).
- Data fusion service that produces VesselFused messages.
- Anomaly detection services that produce AnomalyPrediction message types.
- Sub-detection services that produce AnomalyIntention messages. These services are similar to the anomaly detection services, but aim to detect indications of activities that are not in themselves suspicious, but that could form part of the activities needed to carry out an intention. They can be called weak signals.
- Finally, aggregation services that read both AnomalyPrediction and AnomalyIntention messages and combine them to detect intentions of illegal behaviour, encoded as AnomalyIntention messages.

IV. Services

In AI-ARC, anomaly detection and intent recognition build on the same statistical basis, and they have complementary roles in maritime surveillance domain. Intent recognition allows detection of significant events that are not anomalous, while anomaly detection can detect events that are of interest but not necessarily with intent (such as running on ground, drifting, or operational failures). On a technical level, intent recognition can be defined in terms of anomaly detection by considering intent as following a plan for achieving a goal: an intention is revealed by the detection of multiple indications and anomalies that belong together. In this sense, detection of intent is defined as a sequence of individual anomalies. On a conceptual level, detecting a behaviour or an intent can give meaning to detected anomalies (e.g., the answers to questions like 'why did this ship stop here'). AI-services for maritime surveillance can thus utilize behavioural and intent recognition to further enhance the precision with respect to the anomalies to which attention should be focused. The general approach taken in AI-ARC follows the one presented in [5], [8], [9], while the AI methods generally are based on [10], [11].

A. Sensor data processing services

JDL level 0 and 1 services developed in AI-ARC included detection of unclear waters that could indicate an environmental violation. The system successfully identified both oil spills and algae blooming taking place in the Baltic sea. In addition, satellite images were analyzed to detect ships. This was then correlated with AIS data to find dark vessels.

Figure 2 shows an example of an oil spill detection visualised in the DigLT.

B. Machine learning based anomaly detection services

An important part of level 2 fusion is to find anomalies in the current situation data. AI-ARC developed several services for this.

1) Ship movement prediction: is an essential for anomalous and abnormal behavior of ships. The difficulty of the task grows along the time period of the prediction and the width of the area. This is because there is an exponential divergence typical to nonlinear time series [12], and approximately linear divergence along the area covered. The accuracy and computational load are increased the more data of the past traffic can be accessed by the prediction methods. This is because of the presence of transitions [12] to new dynamics. Although transitions (arriving to archipelago and facing optional routes, reacting to traffic at the present or crossing ship lane) are more restricted than e.g. with a chaotic system, these require data samples to get an adequate coverage for each possible branching of the vessel route. A thematic summary of the reliability (a subjective measure provided to the end-user):

$$\text{reliability} \propto \text{history size}/\text{prediction interval}$$

Meanwhile, the algorithm complexity goes by:

$$\text{complexity} \propto \frac{\text{history size}}{\text{compression ratio}} * \text{prediction interval}$$

The future situational image was decided to have 5 mins update interval during which new predictions should complete. This insight directed us to check many possible data compression methods, which serve as an internal data for the vessel movement prediction.

Experiments with k nearest neighbors (k-NN) method indicate that there is great potential in AIS data compression. The actual vessel movement prediction method provided is based on a constant velocity assumption with a local adaptation. This means the abnormal behavior flag is triggered less eagerly in an area, where there is a turn in a common ship lane (eg.avoiding a peninsula or an island, or turning towards a port).

2) Abnormal patterns in vessel behaviour: are important to discover. A feature engineering of AIS fields plus the analysis of historical big data led to the development of several, configurable machine learning detectors that produce streaming alerts.

The underlying machine learning models are built upon instances of Isolation Forest [13], [14] and Local Outlier Factor [15], [14], which are combined as ensemble classifiers via a majority voting strategy. The maritime area is divided into a customizable grid, yielding a collection of tiles. For each tile, an instance of Isolation Forest and Local Outlier Factor models is trained and fine-tuned using historical AIS data. This approach results in a collection of machine kearning models indexed by the tiles of the grid. This local strategy was preferred over employing a single, global machine learning model for the entire region. These models are stored internally by the service and are encapsulated within a customizable ensemble using a majority-voting strategy. Additionally, there is an orchestration layer that, during the inference stage, selects the corresponding model according to the tile where the prediction is requested.

For every incoming vessel-related message, the machine learning component determines whether it is labeled as abnormal or not. Three abnormal behaviors are within the scope: abnormal speed, abnormal position/trajectory, and a combination of both, referred to as a kinematic anomaly. Additionally, a reliability layer has been implemented to enhance the explainability and robustness of these detections, descried in the next sub-section.

3) Assessing Anomaly Detection Reliability via Feature Importance: A domain-specific reliability assessment (RA) was developed for a Local Outlier Factor-based model [14], [15] designed to detect anomalies in vessel behaviour. The RA's main goal is to improve the interpretability of the model's decision-making process by breaking down the importance of individual features by means of Shapley values [16], thereby making the model more explainable. Anomalies in the AIS data are classified into three categories, according to their cause: speed, position, and kinematic anomalies. Based on this classification, the Most Important Features (MIF) are determined by their respective feature importance. A reliablilty score (RLS) is proposed that combines various metrics, addressing both aspects of robustness and uncertainty, to provide a comprehensive evaluation of the reliability of the model's predictions. The metrics included in the RLS are:

- Anomaly Quantiles: The algorithm's decision function is leveraged to quantify the abnormality of new observations by comparing them to the distribution of anomalies in the training data, using quantiles for assessment.
- Maximal Proportion: The MIF Shapley value of a prediction is related to the maximal shap value observed for that feature in the training data in order to evaluate it's robustness.
- Local Contribution: By normalizing the local feature importance of the MIF, the contribution of a single feature to the prediction is measured, reflecting the confidence level in attributing a prediction to a particular anomaly category, where a higher value indicates increased likelihood.
- Geographic Neighbors: The proportion of the average geographical distance of the nearest neighbors for a given prediction to the maximum average distance of the nearest neighbors among identified anomalies in the training dataset, specifically considering the 'longitude' and 'latitude' features related to the position anomaly category.

The results of these metrics are categorical and combined in the RLS as a mean value, with values ranging between 1 and 5, representing very low and very high reliability, respectively.

C. Statistics based anomaly detection

Other anomaly detection services were based on statistical models in combination with rules appropriate for each case. More specifically, the statistical anomaly detection is based on a framework called "principal anomaly" [11]. The idea is that if a parametric model can be assumed for the observed data (which we will see is possible in many cases), then we

can use Bayesian statistics to produce a predictive distribution for a new observation. If the actual observation has too small probability according to the predictive distribution, then it is considered an anomaly. We now derive the Bayesian principal anomaly. Suppose that we have managed to estimate a, potentially multivariate, probability distribution over the data, $P(x)$. Then we can get a measure of how anomalous a new sample z i, by calculating the probability of getting another sample that is more probable than z:

$$A(z) = \int_{x \in \Omega} P(x) \quad \text{where} \quad \Omega = \{x : P(x) > P(z)\}$$

The value $A(z)$ is between 0 and 1, where 0 means that no samples are more probable than z and close to 1 means that almost all samples are more probable, i.e., to get a sample at least as unusual as z from the distribution of normal data is very unlikely. This could therefore be used directly as a measure of the anomaly. However, since most anomalous values will be very close to 1, a much more convenient measure to use is the negative logarithm of the complement of A:

$$\Lambda(z) = -\log(1 - A(z))$$

This is what we will use as anomaly score. It has values between 0 and plus infinity, where a higher value means more anomalous. By selecting an anomaly threshold, it is possible to control the number of false alarms, epsilon:

$$\Lambda_{\text{anom}} = -\log(\varepsilon)$$

ε is the probability that a normal sample from the distribution will be misclassified as anomalous. It is therefore also the significance level of a hypothesis test of rejecting the null hypothesis that z is not anomalous. It remains to estimate the probability distribution over the data. We will do this by using a Bayesian approach, calculating the predictive distribution of a new sample x given all previous samples D:

$$P(x|D) = \int_{\theta} P(x|\theta)P(\theta|D)$$

Here the integral is over all parameters θ, and the posterior distribution over θ is calculated by Bayes' rule using the prior over θ and the likelihood of the data D:

$$P(\theta|D) \propto P(\theta) \sum_i P(x_i|\theta)$$

For many standard parametric distributions, such as Gaussian, Bernoulli, Gamma, or Poisson, the above integral can be solved analytically, which makes the anomaly calculations very fast. Furthermore, by limiting ourselves to parametric distributions, a robust result can be achieved with much less data compared to more general non-parametric approaches.

The above-described anomaly detection method can be used either in a service by itself, or in combination with a rule-based method in an anomaly detection service, or as input as symptoms in the intention recognition services described below. Specific examples of services where this type of anomaly detection can be used are:

1) Grounding detection.: This service consists of a combination of a statistical anomaly detector for indicating when a vessel is not following a path where vessels of the same draught usually travel, and a rule comparing the draught with the depth ahead of the vessel according to the sea chart. If there are shallow waters ahead, and the vessel is off the ordinary fairway, there will be a grounding risk alarm.

2) Drift detection.: By modelling the movement pattern, in terms of acceleration, deceleration, and turns, vessels that start moving in an irregular way associated with for example drift can be detected.

3) Meeting detection.: A meeting is defined as the close approach of two vessels which are not in a harbour or in vicinity of the coast, at a distance of less than M meters from each other during at least T seconds, where M and T are configurable parameters. During the demonstration, we used M=50m and T=900s. The minimum distance from harbours or the coast is set to 1000m. Not all meetings are suspicious: there are several ship types that are allowed to meet other vessels, such as various authority vessels (coast guard, customs, police, military, etc,...) and certain special purpose or working vessels (pilot, port tender or tugs, supply vessel, rescue vessel, medical vessel, etc). Since the ship type categorization is rather coarse in the AIS standard, currently the detector excludes meetings involving unit operated by administrations or work vessel, assuming that the meeting is not to be considered as suspicious. To be able to check for potential meetings between all vessels in a large geographical area in real time (or faster), the same hexagonal grid is used as in the Grounding service: Each grid cell keeps track of which vessels are within it at any moment, and when checking for vessels in the vicinity of a given vessel, only vessels in the same and neighbouring cells need to be checked.

4) Route deviation.: Deviations from the expected route may indicate a detour to a meeting at sea, or to pick something up at some location. The destination ID field in AIS is (unfortunately) a non-mandatory free text field. It means that often it is not filled in at all or filled in with some non-useful text (e.g. "On a mission" or "Back home again"). However, many larger commercial vessels for goods or passengers make proper use of the field. Still a port may be specified either with its full name (or one of several alternative names and spellings, or occasional misspelling), or as a Location Code (UN/LOCODE) where each port is represented by a two-letter country code and a three-letter location code. To facilitate the analysis (and to pool data corresponding to the same port) we therefore first convert all recognised port names into the corresponding location code. Only destination fields consisting of a single recognized location name or location code are considered by this service. A statistical model is used for assessing what is an "expected" path to the destination: for each destination location code and each hexagonal grid cell on the sea, the speed, course, and transition from previous grid cell is modelled using historical data. This makes up a kind of "vector field" over the sea of the expected movement from each grid cell towards each destination. To assess a new vessel

Fig. 3. Different types of loitering

movement, Bayesian Principal Anomaly is used to calculate whether the path is significantly different from the modelled expected path.

5) Vessel loitering.: Vessels circling over a specific spot, or mapping the seabed, will stand out on the movement pattern detection. This must then be combined with map information of where shipwrecks or sensitive infrastructures are located. Figure 3 shows different types of loitering behaviour.

6) Fishing: Fishing service detects fishing patterns based on AIS data. The service received and handled streaming data via the AI-ARC infrastructure. RISE developed three approaches for providing the service: one rule-based approach, one semi-supervised clustering approach, and one unsupervised clustering approach. The rule-based approach was based on detecting distinctive markers for fishing. First, fishing is typically an activity of fishing vessels; and second, especially for trawling vessels, a characteristic speed. Thirdly, there are characteristic movements indicating fishing. Typical movement trails are, for example, produced while picking up equipment from the sea, heading towards targeted fishing areas (and back to the port of origin), and as well commencing fishing. Other potential data sources we considered were flag status and geolocation of vessels with respect to exclusive economic zones. In the version of the service used during the Baltic demonstration only the vessel type and characteristic fishing vessel speed was used, which was sufficient for detecting most fishing activities in the data under study.

7) Sharp turn: The sharp turn service is designed to detect the kind of sharp turns performed by tanker vessels when cleaning their tanks. Such a tank cleaning turn is characterized by rapidly turning at a high speed to achieve a high centripetal force, possibly in a pulsating manner, making the vessel rock from side to side and causing water at the bottom of the otherwise emptied tank to flow back and forth through it. The water is then flushed out in the sea. The manoeuvre is usually repeated at least once, with some 10-20 minutes interval. Cleaning the tanks in this way is allowed for certain chemicals and at certain locations, but if it is done at the wrong place or leaves traces at the surface, then it is an envronmental regulation violation.

The service works by calculating the centripetal acceleration of tanker vessels, and checking whether it supersedes a threshold. The centripetal acceleration is calculated from a pair of AIS messages (consecutive but at least 3 seconds apart, to limit noise) providing its rate of turn (RO_{TAIS}) by multiplying the mean speed of the two readings with the difference in course (in radians) and dividing by the difference in time.

D. Intention services developed

JDL level 3 services are intended for impact or threat assessment. Within the scope of AI-ARC, this means that we must detect possible criminal intentions of vessels.

What is common for the intention recognition methods is that they try to identify a plan, goal, or intent, behind some observed behaviours. That is, rather than just observing the current speed and position of a vessel, and trying to predict where it will be some time later, the question posed is *why* it has this speed and position, i.e., inferring the longer-term plan. That way it is possible to predict the next steps in the plan or to recognize the goal of the plan.

By prototype-based we mean that several situations of interest to detect have been identified as "use cases" in the project, and these situations will be characterized as "prototypes," describing typical appearances of the situations. Such prototypes can then be the basis for classification or detection, using real data as input. Prototypes can either be specified manually by domain experts, trained from example data, or be initially specified by experts and then enhanced from data over time. The prototype-based approach for intention recognition is inspired by the threat analysis methodology used by intelligence agencies [9] to combine weak signals: Try to identify potential threats to different vulnerabilities by describing how an attack could be performed, and then look for evidence in data of someone performing those steps.

1) The Environment service: The environment service is an aggregating service to detect environmental regulation violation. The use case we have focused on in this project is cleaning of tanks in a prohibited way. Cleaning of tanks is allowed for some chemicals, and if it is done at certain locations (with sufficient depth) and without leaving traces on the surface.

The service combines input from the Sharp turn service, to detect ships that perform characteristic tank cleaning movements), the Dark vessel service, to detect ships that for some reason are shutting down their AIS), and the Oil spill detection service, to detect visible traces of oil or other chemicals on the surface). See Figure 4 for an overview. The combination of inputs is rule based: When an oil spill is detected in some region, any previously detected dark vessel, or (non-dark) vessel detected as performing a number of tank cleaning movements, within a given radius from the oil spill and up till 12 hours prior to the oil spill detection, are returned as possible sources of the oil spill.

2) Smuggling: The smuggling service is an aggregating service to detect behaviours that may be considered as abnormal and may be associated with smuggling activities. The specific type of smuggling we have focused on in this project is that a larger vessel (assumed to have AIS with regards to International Maritime Organisation regulation) meets up with a smaller vessel somewhere (potentially not using its AIS) to hand over some goods in either direction, alternatively that the goods is dropped in the sea from one vessel, to be picked up later by the other vessel. The service works by combining weak signals in the form of Anomaly Indication messages

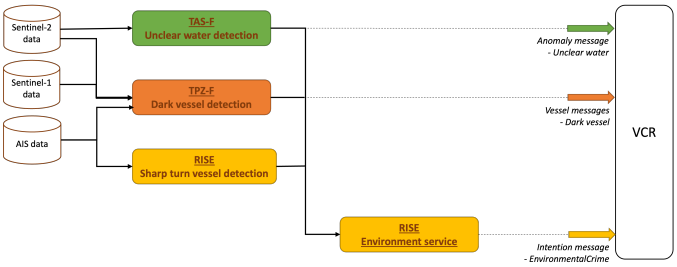

Fig. 4. Service composition for the environment case

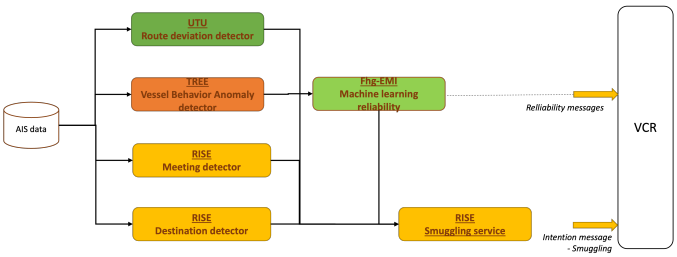

Fig. 5. Smuggling service composition.

from other services in the chain, and issuing an alarm when the combined evidence is considered strong enough. For the aggregation, the service uses a variant of the Prototype-based intent recognition technique described above.

The service composition used in the smuggling use case is shown in Figure 5.

While it is not possible to confidently detect smuggling solely based on AIS positional data, we saw that it is possible to detect interesting features related to smuggling and other crimes. Both presented cases involved meetings with small non-AIS vessels. With radar coverage, and future short wavelength radar, such vessels might be detectable. Both presented cases involved seemingly unmotivated detours (but the 2010 one was clever enough to adjust the destination accordingly). The vessels move ”impeccable” during handover – no slow down or unsteady path. However, when intercepted by two vessels and forced to change course all our detectors went off – reminiscent of other potential use case: Piracy

3) Dynamic Bayesian Networks: The final aggregating service used was based on Dynamic Bayesian Networks [17].

The Dynamic Bayesian Network (DBN) editor provides the end-user with the possibility to create, edit and validate DBNs.

The evidence generator gives on the one hand the possibility to check for specific conditions by using a simple logic-based language, e.g., checking “if the speed of a vessel is higher than 10kts and the vessel is in a specific area”, and on the other hand to test the evidence generation with test data. For each node in the DBN, the end-user is then able to connect the random variable it represents to evidence from the evidence generator. After building the DBN, the end-user can validate the functionality of the network by setting evidence and perform the inference to generate the probabilities for each random variable.

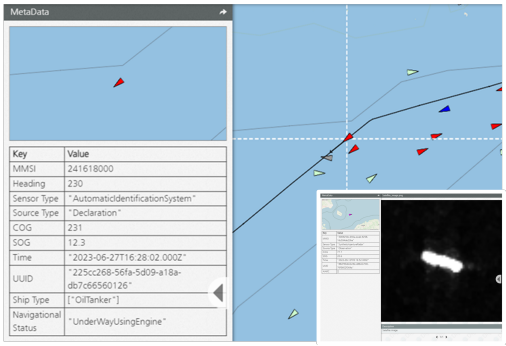

Fig. 6. DigLT example.

In the first use case, a service chain was created to detect illegal fishing. The use case was divided into two subtasks, detection of fishing movements in (time) restricted areas and reporting of dark vessels in restricted areas. Included as input was a fishing activity detector by and a dark vessel detector.

The objective of the critical infrastructure threat service was to identify suspicious behaviour around maritime critical infrastructure, such as submerged pipelines or cables.

E. Risk index computation

This unified risk index computation approach is primarily driven by the implementation of fuzzy systems methodology integrated with Multiple Criteria Decision Making (MCDM) methodology for prioritisation and sensitivity analysis.

The risk index computation Service was used in the grounding use-case, which assesses the risk of vessels running aground.

V. Visualization and User Interaction

Visualization and user interaction was provided primarily by the DigLT system, as shown in Figure 6 above. The system also provides the opportunity for 3D visualisation in virtual reality, but this was not used in the Baltic demonstration.

VI. Validation and Evaluation of the Baltic AI-ARC demo

The Baltic demo was evaluated by questionnaires to invited end-users and experts. In summary, the results were favourable, but also indicated that some more research is needed. A set of Key Performance Indicators (KPI’s) were developed early in AI-ARC, and the services shown in the Baltic demo managed to meet 90 % of these. In addition to the questionnaires filled in by law enforcement stakeholders, also civilian mariners were interviewed. They highlighted the potential of the VCR as a support to navigation in difficult conditions.

The questionnaire consisted of some Likert scale 1 to 6 questions and open questions. A summary of the Likert scale question results is shown in Table I.

The special observation from the platform questions is that the last question, “The system presented would provide added

TABLE I
OVERALL RATINGS OF THE DIFFERENT VCR SERVICES

	Objective Met	Increased Awareness	Useful Output	Clear Outcome
Average all	4.79	4.78	4.78	4.88
Operational partners	4.9	4.95	5.15	5.1
Technical partners	4.85	5.02	4.93	5.18
External	3.95	3.95	4.00	3.85

value for my organisation", was evaluated by the operational partners to rate 5.75.

The evaluation also included a qualitative part. The most important observations in this were:

- For anomaly detection the combination of radar and satellite data would give remarkable added value
- Symbolics should be further developed. Critical Alerts should be "flagged" as a top priority clearly, sound combined with alerting
- AI-ARC shortens the reaction or response times
- Great added value since it detects remotely and could strengthen court cases (Illegal fishing, external)
- Automated flagging feature in this service is valuable for protecting CI (Critical Infrastructure, operational partner)
- Wonder why this type of service is not in use? Super much added value when protecting for our maritime nature and improvement to EMSA service is very high (Environmental violation, operational partner)
- Maybe a policy initiative should be established for within the Europea Commission's Directorate-General for Migration and Home Affairs (DG HOME) or Directorate-General for Maritime Affairs and Fisheries (DG MARE) to make this service toolbox happen! (Any proposals for revisions and/or additions to the requirements and specifications, operational partner)

VII. SUMMARY AND FUTURE WORK

In summary, we presented an overview of the level 1, 2, and 3 fusion services implemented in the AI-ARC maritime surveillance system of systems. Use cases and evaluation results from a demonstration in the Baltic sea focusing on detecting illegal activities were described. The evaluation shows that a micro-service based solution, with services that are composable and can be adapted by the users for different use cases, along with the DigLT visualization platform can contribute significantly to increased maritime situation awareness.

The AI-ARC project has also developed a road map for future deployment of the solution, and are exploring the possibilities to use the approach also in other domains. A follow-up project will focus on protection of critical underwater infrastructure.

ACKNOWLEDGMENT

We thank all our colleagues in the AI-ARC project team for fruitful discussions and pleasant collaboration.

REFERENCES

[1] S. F. Andler, M. Fredin, P. M. Gustavsson, J. van Laere, M. Nilsson, and P. Svenson, "SMARTracIn: a concept for spoof resistant tracking of vessels and detection of adverse intentions," in *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense VIII* (E. M. Carapezza, ed.), vol. 7305, p. 73050G, International Society for Optics and Photonics, SPIE, 2009.

[2] M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," *Human Factors*, vol. 37, no. 1, pp. 32–64, 1995.

[3] B. C. Dostal, "Enhancing situational understanding through the employment of unmanned aerial vehicles." Interim Brigade Combat Team Newsletter, 2007.

[4] R. Popp and J. Poindexter, "Countering terrorism through information and privacy protection technologies," *IEEE Security & Privacy*, vol. 4, pp. 18–27, Nov.-Dec. 2006.

[5] J. Brynielsson, A. Horndahl, L. Kaati, C. Mårtenson, and P. Svenson, "Development of computerized support tools for intelligence work," in *Proceedings of the 14th International Command and Control Research and Technology Symposium (14th ICCRTS)*, (Washington, DC), p. 48, 2009.

[6] M. E. Liggins, D. L. Hall, and J. Llinas, *Handbook of multisensor data fusion: theory and practice; 2nd ed.* Electrical engineering and applied signal processing series, Hoboken, NJ: Taylor & Francis Ltd, 2008.

[7] E. P. Blasch and S. Plano, "Level 5: user refinement to aid the fusion process," in *Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2003* (B. V. Dasarathy, ed.), vol. 5099, pp. 288 – 297, International Society for Optics and Photonics, SPIE, 2003.

[8] P. Svenson, R. Forsgren, B. Kylesten, P. Berggren, W. Fah, M. Choo, and J. Hann, "Swedish-singapore studies of bayesian modelling techniques for tactical intelligence analysis," in *2010 13th International Conference on Information Fusion*, pp. 1–8, 2010.

[9] J. Brynielsson, A. Horndahl, F. Johansson, L. Kaati, C. Mårtenson, and P. Svenson, "Harvesting and analysis of weak signals for detecting lone wolf terrorists," *Security Informatics*, vol. 2, 2013.

[10] D. Gillblad, R. Steinert, and A. Holst, "Fault-tolerant incremental diagnosis with limited historical data," in *International Conference on Prognostics and Health Management 2008 (PHM'08)*, 2008.

[11] A. Holst and J. Ekman, "Incremental stream clustering for anomaly detection and classification," in *Eleventh Scandinavian Conference on Artificial Intelligence*, IOS Press, 2011.

[12] B. Goswami, "A brief introduction to nonlinear time series analysis and recurrence plots," *Vibration*, vol. 2, pp. 332–368, 2019.

[13] F. Liu, K. Ting, and Z. Zhou, "Isolation forest," in *Proceedings of the 8th IEEE International Conference on Data Mining*, pp. 413–422, IEEE, 2008.

[14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[15] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "Lof: identifying density-based local outliers," in *Association for Computing Machinery*, vol. 29, (New York, NY, USA), pp. 93–104, June 2000.

[16] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30*, pp. 4765–4774, 2017.

[17] P. Dagum, A. Galper, and E. Horvitz, "Dynamic network models for forecasting," in *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*, pp. 41–48, AUAI Press, 1992.