Dark-VADER: Detection of Anomalous AIS Message Delays for Maritime Situational Awareness

Giorgio Ioannou[†*], Domenico Gaglione[*], Leonardo M. Millefiori[*], Alfredo Renga[†], Paolo Braca[*], Peter Willett[‡]
[†]University of Naples Federico II, Naples, Italy
[*]NATO STO Centre for Maritime Research and Experimentation, La Spezia, Italy
[‡]University of Connecticut, Storrs, CT, USA

***Abstract*—Maritime situational awareness (MSA) refers to the effective understanding of activities related to maritime environment. Central to MSA, particularly concerning non-military vessels, is the automatic identification system (AIS), which provides real-time data on vessel movements. However, anomalies such as intentional AIS transponder disablement pose significant challenges to MSA, potentially indicating illicit activities. This paper introduces the Dark-VADER (dark vessel AIS delay event recognition) algorithm, designed to detect AIS switch-offs by comparing the frequency of message reception from a vessel under examination with that of neighboring vessels. Leveraging a statistical hypothesis testing procedure based on a Bernoulli process, the algorithm distinguishes between normal and anomalous behavior. Validation using real-world AIS data confirms the fitness of the selected distribution model for times between message arrivals, essential for the algorithm's operation. Overall, this preliminary work provides a foundational framework for improving maritime AIS anomaly detection, with avenues for future development towards more robust and dynamic approaches.**

***Index Terms*—Anomaly detection, data analysis, parameter estimation, marine technology, maritime situational awareness**

I. Introduction

Maritime situational awareness (MSA), i.e., the capability to effectively understand all the activities related to the maritime environment, predict their evolution, and assess their level of risk [1]–[3], strongly relies on the data provided by the automatic identification system (AIS) [4]–[6], especially in the context of the "white picture," related to non-military assets.

Originally conceived only for safety and collision avoidance, the AIS represents today the largest source of information on maritime vessel traffic. Indeed, the International Maritime Organization (IMO) mandates [7] that AIS is compulsorily operated on: ships over 300 gross tonnage (GT), cargo vessels over 500 GT, all passenger ships, and all fishing vessels over 45 meters (over 15 meters in EU countries [8]). The AIS transponder is used to broadcast messages containing both the vessel's dynamic (e.g., position, velocity, course) and static (e.g., identifier, dimension) information over two very high frequency (VHF) channels, which are then received by neighboring vessels, as well as by terrestrial and satellite receivers.

This work was supported by the NATO Allied Command Transformation (ACT) under the Data Knowledge and Operational Effectiveness (DKOE) project.

Messages are sent at regular intervals, and to use efficiently the available bandwidth, vessels that are anchored or slowly moving transmit less frequently than those that are navigating faster or maneuvering. Nonetheless, transmitted messages might not be received due to a congested channel or degraded propagation conditions, or as a result of an intentional or accidental disablement of the transponder, i.e., a switch-off. The AIS specification does not establish any acknowledgment mechanism for well received messages, thus meaning that lost messages are not transmitted again; this makes the communication channel *memoryless*.

The very large volume of AIS messages broadcast by vessels all around the globe and the potential use of this technology for security purposes have motivated many researches on AIS anomaly detection algorithms. Among the several identified anomalies [9], the switch-off of an AIS transponder is of high interest as possibly related to an illegal activity that the vessel is conducting whilst concealed [10]. The trajectories of the ships involved in these activities tend to deviate from standard routes. Thus, in different works the AIS disablement problem is addressed by adopting a proper kinematic motion model of the observed vessel [11]–[13], or through the analysis of the entire AIS track to detect anomalous events [14]–[16].

Although the switch-off could be hypothesized from the persistent missed reception of AIS messages, the absence of messages might also be ascribed to poor communication conditions. For this reason, authors in [17] proposed an algorithm for the detection of AIS transponder switch-offs based on a hidden Markov model that characterizes both the state of the onboard transponder and the state of the VHF channels. Similarly, the method described in [18] leverages the received signal strength indicator to detect whether the absence of AIS messages is representative of an anomalous event, e.g., the switch-off of the transponder, or it is due to a communication dropout. The work presented in [19], instead, does not focus directly on the detection of the transponder switch-off as such, but it addresses the problem of recognizing a route deviation that may occur during the disablement of the AIS transponder and that might conceal an illegal activity. Finally, an anomaly detection framework that classifies intentional and non-intentional AIS switch-offs based on an artificial neural network (ANN) is proposed in [20].

In this article, we propose Dark-VADER (dark vessel AIS delay event recognition), an algorithm for the detection of a

vessel's AIS transponder switch-off based on both the absence of messages received from the vessel itself and the messages correctly received from the neighboring vessels. Indeed, the number and rate of messages received from all the vessels in a given area can be seen as representative of the quality of the communication channel, and this information can decrease the chances of erroneously detecting a switch-off when the absence of messages is rather due to poor communication conditions. The detection problem is formalized as a hypothesis testing procedure. The successful or unsuccessful reception of AIS messages from the vessel under test is modeled as a Bernoulli process with a given probability of success, that is approximately the rate of received messages. Roughly speaking, comparing this rate with that of the neighboring vessels, the hypothesis testing procedure decides between the normal or anomalous behavior of the vessel under test.

The statistical model, based on the Bernoulli and geometric process assumptions, is verified on real-world AIS data from a terrestrial receiver operated in the area of La Spezia, Italy, by the NATO Science and Technology Organisation (STO) Centre for Maritime Research and Experimentation (CMRE). The area and dataset used for this analysis are shown in Fig 1. The remainder of this article is organized as follows. Section II introduces the dark vessel detection problem and provides a model for the AIS message arrivals. The Dark-VADER decision statistic and its theoretical performance are derived in Sec. III. Section IV presents the validation of the geometric process assumption, while the empirical performance of Dark-VADER and some non-idealities are discussed in Sec. V. Concluding remarks are provided in Sec. VI.

II. Problem Formulation

Let us consider a set of J vessels, indexed by $j = 1, \dots, J$ in a region of interest, which report their kinematic information via AIS messages. Such messages are normally broadcast at regular intervals according to the AIS specification [21]. More concretely, the reporting interval at a given time t, namely $\Delta_j(t)$, can take values from a finite set, depending on the vessel's kinematics, and with finer precision on its speed and course over ground. In our setup, the AIS messages are received by a single receiving base station, either coastal, ship- or space-borne. The time interval between two subsequent received messages, namely the *interarrival time*, is not necessarily equal to $\Delta_j(t)$, due to possible communication errors, interference, channel congestion, etc. Consequently, the sequence of messages received from vessel j arrive at time instants $t_{j,i}$, with $i = 0, \dots, N_j$, where $N_j + 1$ is the total number of received messages. These time instants identify N_j intervals that are spaced in time by a quantity proportional to $\Delta_j(t)$:

$$t_{j,i+1} - t_{j,i} = \delta_{j,i+1} \cdot \Delta_j(t), \quad \text{with } \delta_{j,i+1} \in \mathbb{N}, \tag{1}$$

where $t \in (t_{j,i}, t_{j,i+1}]$, and in this time interval $\Delta_j(t)$ is constant. The value of $\Delta_j(t)$ is assumed *known* at the receiver side, since it can be deduced from the vessel's kinematics.

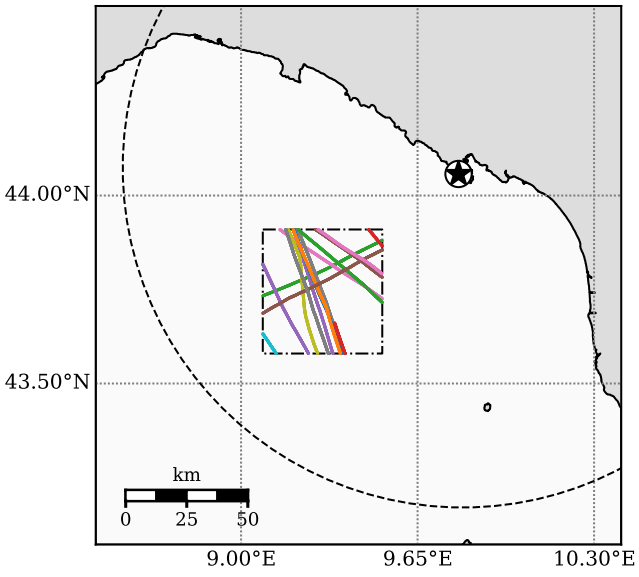

Fig. 1. Real-world dataset of AIS messages in a selected area off the Gulf of La Spezia. The dataset comprises tracks from 33 vessels in a 24 h period between 4 and 5 July, 2023. Each color identifies a different vessel. The position of the AIS receiver is marked by a black circled star. The dashed circle represents a 100 km radius around the receiver.

Consequently, we can map the time instants of the message arrivals on a discrete axis, ruled by:

$$\delta_{j,i} = \frac{t_{j,i} - t_{j,i-1}}{\Delta_j(t)}, \quad i = 1, \dots, N_j. \tag{2}$$

From a modeling standpoint, the key innovation of this work is to represent the discrete intervals $\delta_{j,i}$ as random variables geometrically distributed. The geometric assumption will be validated against real-world data in the following, but it is also well grounded from a theoretical perspective given that the AIS communication protocol (including both the transmission and the communication) is memoryless.

We assume that discrete intervals from vessels that are close in space and time must follow the same geometric distribution. This is a fundamental assumption to detect the intentional or unintentional disablement of a vessel's AIS transponder. More formally, we propose a statistical hypothesis testing procedure, named Dark-VADER, whose null hypothesis, denoted by $\mathcal{H}_0$, represents the normal operation of the AIS transponder of a vessel under test $j_{\mathrm{T}} \in \{1, 2, \dots, J\}$. The alternative hypothesis, denoted by $\mathcal{H}_1$, represents instead a switch-off of the AIS transponder onboard the vessel under test.

III. Dark-VADER Statistical Hypothesis Testing

A. Model assumption

Consider a vessel under test $j_{\mathrm{T}} \in \{1, 2, \dots, J\}$ and the sequence of discrete intervals $\mathcal{T}_{j_{\mathrm{T}}} = \{\delta_{j_{\mathrm{T}},i}\}_{i=1}^{N_{j_{\mathrm{T}}}}$, which are independent and identically distributed (i.i.d.) random variables, with geometric distribution of parameter $p_{j_{\mathrm{T}}}$, i.e., $\delta_{j_{\mathrm{T}},i} \sim \mathcal{G}(p_{j_{\mathrm{T}}})$.

It is convenient to consider the analogous Bernoulli process $X_{j_T,k}$, which can be defined as follows:

$$X_{j_T,k} = \begin{cases} 1 & k \in \mathcal{K}_{j_T}, \\ 0 & \text{otherwise}, \end{cases} \tag{3}$$

where $\mathcal{K}_{j_T} = \left\{ \sum_{i'=1}^{i} \delta_{j_T,i'} \right\}_{i=1}^{N_{j_T}}$. Clearly, the probability of success is equal to p_{j_T}, that is the parameter of the discrete intervals geometric distribution, i.e., $\mathcal{P}(X_{j_T,k} = 1) = p_{j_T}$.

Let us denote with $\mathcal{B}(p)$ the Bernoulli distribution of parameter p, and with p^* the nominal rate of messages correctly received under normal circumstances; this parameter can be estimated from the vessels close to the one under test. Neglecting the dependence on j_T from now on to keep the notation light, the statistical hypothesis test is formulated as follows:

$$\begin{cases} \mathcal{H}_0: & X_k \sim \mathcal{B}(p_0), \quad p_0 = p^*, \\ \mathcal{H}_1: & X_k \sim \mathcal{B}(p_1), \quad p_1 < p^*, \end{cases} \tag{4}$$

for $k = 1, \ldots, K$, with $K = \max(\mathcal{K})$. The parameter p_1 under $\mathcal{H}_1$ represents a rate of received messages lower than the nominal parameter p^*. To give an idea, if the vessel intentionally goes AIS dark, i.e., it turns off the AIS transponder, in an area with no communication issues, we expect the parameter p_1 to be significantly lower than the nominal value p^*. The test performance is defined in terms of false alarm probability P_{FA}, i.e., the probability that the test declares the hypothesis $\mathcal{H}_1$ under $\mathcal{H}_0$, and detection probability P_D, i.e., the probability that the test declares $\mathcal{H}_1$ under $\mathcal{H}_1$, meaning in the presence of an actual switch-off of the AIS transponder onboard the vessel under test.

We adopt the Neyman-Pearson (NP) criterion, which maximizes P_D for a given level of P_{FA} under a simple hypothesis testing problem (see, e.g., [22, Th. 3.1]). We observe, however, that in our formulation (4) the distribution under $\mathcal{H}_1$ is not completely defined, because the parameter p_1 is unknown. We therefore adopt the NP criterion assuming p_1 known, and show that both the decision statistic and P_{FA} do not depend on p_1. The NP criterion is based on the use of the *likelihood ratio* (LR) as decision statistic, which in our case particularizes as

$$\begin{aligned} \mathcal{L} &= \frac{\mathcal{P}(X_1, \ldots, X_K; \mathcal{H}_1)}{\mathcal{P}(X_1, \ldots, X_K; \mathcal{H}_0)} \\ &= \frac{\prod_{i=1}^{K} p_1^{X_i}(1-p_1)^{1-X_i}}{\prod_{i=1}^{K} p_0^{X_i}(1-p_0)^{1-X_i}} \\ &= \frac{p_1^{\sum_{i=1}^{K} X_i}(1-p_1)^{\sum_{i=1}^{K}(1-X_i)}}{p_0^{\sum_{i=1}^{K} X_i}(1-p_0)^{\sum_{i=1}^{K}(1-X_i)}} \\ &= \frac{p_1^{S_K}(1-p_1)^{K-S_K}}{p_0^{S_K}(1-p_0)^{K-S_K}}, \end{aligned} \tag{5}$$

where in the last equality we have defined the random variable $S_K = \sum_{i=1}^{K} X_i$ as the number of successfully received AIS messages. The test declares $\mathcal{H}_1$ if the LR in (5) is larger than a threshold γ', which is set after a given false alarm level, $P_{FA}(\gamma')$. As usual in several contexts, it is possible to further simplify the decision statistic, taking the logarithm of the LR, i.e., $\ell = \log \mathcal{L}$, and obtaining an equivalent test based on the log-likelihood ratio (LLR):

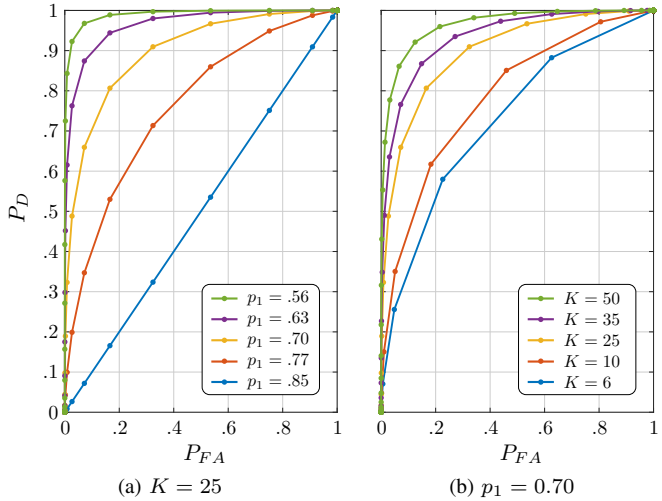

(a) $K = 25$ (b) $p_1 = 0.70$

Fig. 2. Receiver operating characteristics for different values of the parameters p_1 and K, in the case of a nominal rate of received message $p^* = 0.85$.

$$\begin{cases} \ell \le \log \gamma': & \text{decide } \mathcal{H}_0, \\ \ell > \log \gamma': & \text{decide } \mathcal{H}_1. \end{cases} \tag{6}$$

The LLR takes the following form:

$$\begin{aligned} \ell &= \log \left[\left(\frac{p_1}{p_0} \right)^{S_K} \left(\frac{1-p_1}{1-p_0} \right)^{K-S_K} \right] \\ &= S_K \log \left(\frac{p_1}{p_0} \right) + (K - S_K) \log \left(\frac{1-p_1}{1-p_0} \right) \\ &= S_K \left[\log \left(\frac{p_1}{p_0} \right) - \log \left(\frac{1-p_1}{1-p_0} \right) \right] + K \log \left(\frac{1-p_1}{1-p_0} \right). \end{aligned} \tag{7}$$

Since $p_1 < p_0$, the multiplicative coefficient of S_K in (7) is always negative. Consequently, the statistical hypothesis test (6) is equivalent to

$$\begin{cases} S_K > \gamma: & \text{decide } \mathcal{H}_0, \\ S_K \le \gamma: & \text{decide } \mathcal{H}_1, \end{cases} \tag{8}$$

where we have defined[1]

$$\gamma := \frac{\log \gamma' - K \log \left(\frac{1-p_1}{1-p_0} \right)}{\log \left(\frac{p_1}{p_0} \right) - \log \left(\frac{1-p_1}{1-p_0} \right)}. \tag{9}$$

The statistical hypothesis testing (8), as well as the procedures to set the threshold and the false alarm probability, and to estimate the nominal parameter p^*, represent the core of the proposed Dark-VADER algorithm.

B. Theoretical performance

The decision statistic is the number of messages successfully received, namely the number S_K of successes in K trials.

[1]It is clear that the expression (9) is not useful to set the false alarm level P_{FA}, as usual in a problem of detection.

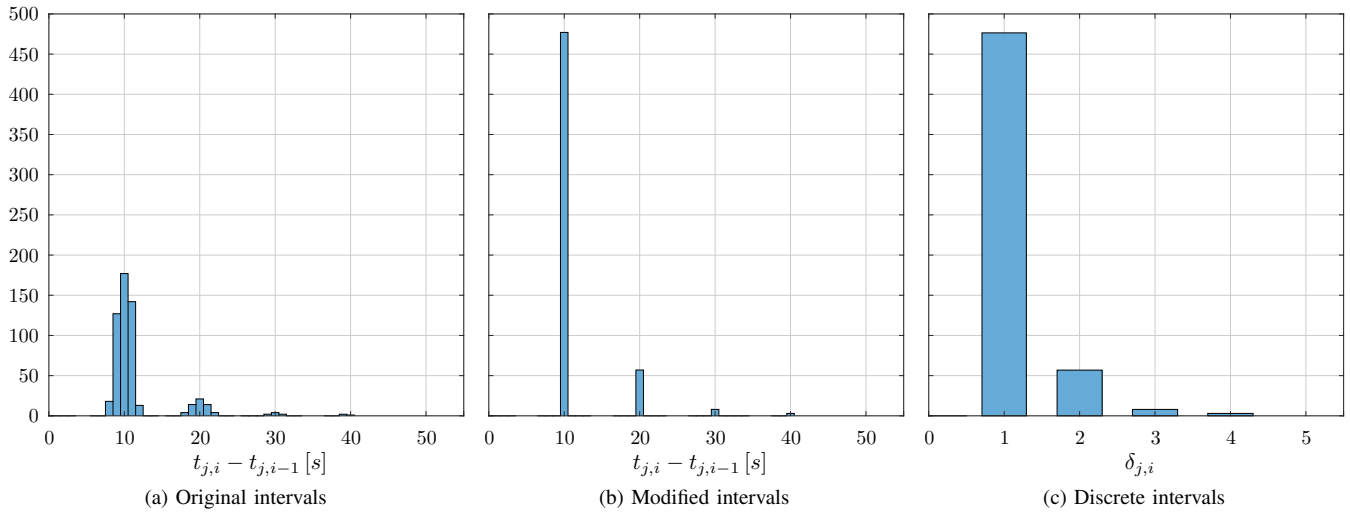

(a) Original intervals (b) Modified intervals (c) Discrete intervals

Fig. 3. Data from vessel j when its reported SOG is less than or equal to $14\,\mathrm{kn}$, i.e., when its reporting intervals is $10\,\mathrm{s}$. (a) Original intervals, recorded by the receiver. (b) Modified intervals, obtained as described in Sec. IV. (c) Discrete intervals, obtained according to (2).

Denoting with $\mathcal{B}(n,p)$ the binomial distribution with n number of trials and probability of success p, the decision statistic S_K is distributed as follows:

$$\begin{cases} \mathcal{H}_0: & S_K \sim \mathcal{B}(K, p_0), \quad p_0 = p^*, \\ \mathcal{H}_1: & S_K \sim \mathcal{B}(K, p_1), \quad p_1 < p^*. \end{cases} \tag{10}$$

Based on the assumptions made at the beginning of this section, it is possible to provide analytical expressions for the receiver operating characteristic (ROC) curve, i.e., for the false alarm and detection probabilities, which are

$$\begin{aligned} P_{FA}(\gamma) &= \mathcal{P}(S_K \leq \gamma; \mathcal{H}_0) = F_{\mathcal{B}}(\gamma; K, p^*), \\ P_D(\gamma) &= \mathcal{P}(S_K \leq \gamma; \mathcal{H}_1) = F_{\mathcal{B}}(\gamma; K, p_1), \end{aligned} \tag{11}$$

where $F_{\mathcal{B}}(s; n, p)$ refers to the cumulative distribution function (CDF) of a binomial of parameters n and p. The two plots in Fig. 2 show the ROC curves for different values of the parameters p_1 and K, and for p^* set to 0.85.

Remark A. It is worth noting that the test (8) no longer depends on the parameters p^* and p_1, which in principle could be unknown, especially p_1 that is related to the anomalous behaviour of the vessel under test. On the other hand, the knowledge of p^* is required to set the desired false alarm level P_{FA}. In the next section we will show how it is possible to set the false alarm level with real-world data.

Remark B. The normalized decision statistic $\frac{S_K}{K}$ can be intuitively interpreted as an estimator $\hat{p}$ of the parameter p, equal to p^* under $\mathcal{H}_0$ and to p_1 under $\mathcal{H}_1$. It represents how often we receive messages from the vessel under test. If $\hat{p}$ is significantly lower than a nominal value, say $p_\gamma = \frac{\gamma}{K}$ [see (10)], then the vessel is not transmitting often enough its messages and we declare $\mathcal{H}_1$. Otherwise, if $\hat{p}$ is larger than this nominal value, we have enough evidence that the vessel is broadcasting AIS messages normally.

IV. Preprocessing and Model Validation

In this section we validate the proposed model on real-world data recorded by an AIS terrestrial receiver operated by CMRE in the area of La Spezia, Italy. The main advantage of using the data of a single AIS receiver rather than from a network of AIS receivers is that we can trust the dataset is not decimated, which allows us to verify that the temporal spacing of AIS arrivals is actually ruled by (1). Nonetheless, Dark-VADER can be applied to decimated AIS data, too, assuming that the decimation strategy is known.

We have considered a circular area of radius $100\,\mathrm{km}$ from the AIS receiver, as shown in Fig 1, in which an average of $800\,000$ position reports are recorded every day. Each position report message contains fifteen fields, among which only the following are needed for our analysis: maritime mobile service identity (used to identify a vessel and extract its trajectory), latitude and longitude, course over ground (COG), speed over ground (SOG), and navigational status (used to distinguish whether the vessel is moored or under way). Besides these fields, we clearly use also the timestamp of the message; this information, however, is not included in the message itself, but recorded by the base station upon reception.

As stated in Sec. II, the reporting interval $\Delta_j(t)$ depends on the vessel's kinematics. More specifically, according to the AIS regulation [21] the reporting interval is set to $10\,\mathrm{s}$ for a vessel navigating with SOG below $14\,\mathrm{kn}$, to $6\,\mathrm{s}$ for a vessel navigating between $14\,\mathrm{kn}$ and $23\,\mathrm{kn}$, and to $2\,\mathrm{s}$ for a vessel navigating faster than $23\,\mathrm{kn}$. This rule is more complex when the COG varies over time, i.e., when the vessel is maneuvering. To facilitate our analysis and take a step-by-step approach, we have chosen to focus solely on non-maneuvering vessels in this study, leaving to future works the validation against COG-changing vessel data.[2] To this aim, we select an area

[2]Note that Dark-VADER can easily deal with AIS data from maneuvering vessels, too, as long as the reporting rules are clearly identified.

TABLE I
VALIDATION OF THE GEOMETRIC DISTRIBUTION MODEL: RESULTS FOR SOME VESSELS THAT NAVIGATES AT DIFFERENT SOGS.

	SOG $< 14\,\text{kn}$			$14\,\text{kn} \leq$ SOG $< 23\,\text{kn}$			$23\,\text{kn} \leq$ SOG		
j	N_j	$\hat{p}_j$	RMSE_j	N_j	$\hat{p}_j$	RMSE_j	N_j	$\hat{p}_j$	RMSE_j
5	-	-	-	573	0.907	0.003	963	0.871	0.012
6	-	-	-	-	-	-	1340	0.941	0.003
10	-	-	-	949	0.749	0.018	-	-	-
15	545	0.869	0.006	819	0.881	0.007	-	-	-
24	-	-	-	587	0.854	0.012	-	-	-
30	421	0.822	0.007	30	0.857	0.017	59	0.776	0.039
31	433	0.738	0.010	-	-	-	-	-	-

With a small abuse of notation, N_j, $\hat{p}_j$ and RMSE_j in this table refer to quantities computed not on the entire set of messages from vessel j, but on the SOG-conditioned message subset indicated in the header.

in the Ligurian Sea off the Gulf of La Spezia, where vessels typically do not maneuver. This area, along with the real-world dataset used for validation, is illustrated in Fig. 1.

As mentioned, the messages received from the j-th vessel include, among other information, the SOG, a time-varying value that we refer to as $\text{SOG}_j(t)$. Depending on $\text{SOG}_j(t)$, the reporting interval $\Delta_j(t)$ takes on the values 10 s, 6 s, and 2 s. In order to compute the discrete intervals $\delta_{j,i}$ according to (2) and validate the geometric distribution assumption, we calculate the interarrival times $t_{j,i} - t_{j,i-1}$, for $i = 1, \ldots, N_j$, and group them according to the value of SOG at time $t_{j,i}$, i.e., $\text{SOG}_j(t_{j,i})$. In other words, if $\text{SOG}_j(t_{j,i}) < 14\,\text{kn}$, the i-th interarrival time $t_{j,i} - t_{j,i-1}$ is associated to the longest reporting interval, i.e., $\Delta_j(t_{j,i}) = 10\,\text{s}$. Likewise, if $14\,\text{kn} \leq \text{SOG}_j(t_{j,i}) < 23\,\text{kn}$, the i-th interarrival time is associated to the reporting interval $\Delta_j(t_{j,i}) = 6\,\text{s}$, whereas if $\text{SOG}_j(t_{j,i}) \geq 23\,\text{kn}$, then the i-th interarrival time is associated to the shortest reporting interval $\Delta_j(t_{j,i}) = 2\,\text{s}$. As an example, Fig. 3a shows the histogram of the interarrival times computed for a given vessel j and associated to the longest reporting interval of 10 s. As expected, the histogram presents multiple modes, the first one located at 10 s, and the others located at multiples of 10 s. These additional modes are due to single messages or sequence of messages from the vessel that were not received by the base station. Moreover, there are non-zero occurrences of interarrival times that are different from 10 s or its multiples, but that are next to the modes; these are caused by some non-idealities of the AIS reception system. To deal with these non-idealities, the interarrival times different from 10 s and its multiples are replaced with the closest multiple of 10 s; this leads to the modified histogram in Fig. 3b. Finally, according to (2), these interarrival times are divided by the corresponding reporting interval, i.e., 10 s, so to obtain the discrete intervals $\delta_{j,i}$, whose histogram is reported in Fig. 3c. We note that this two-step processing aimed at dealing with the non-idealities of the AIS communication protocol is equivalent to computing the discrete intervals as:

$$\delta_{j,i} = \text{round}\left(\frac{t_{j,i} - t_{j,i-1}}{\Delta_j(t_{j,i})}\right)$$

rather than as in (2), where $\text{round}(\cdot)$ approximates the argument to the nearest integer.

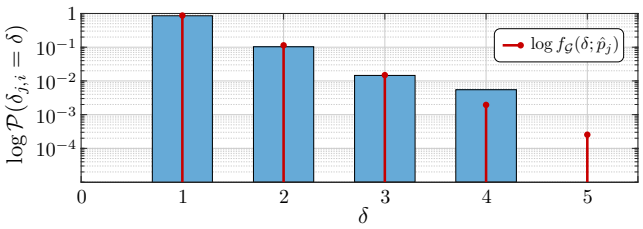

(a) Vessel $j = 15$ at SOG $< 14\,\text{kn}$ ($\hat{p}_j = 0.869$)

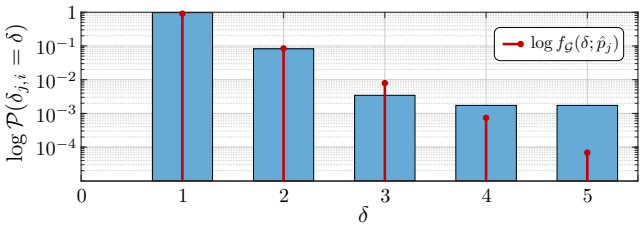

(b) Vessel $j = 5$ at $14\,\text{kn} \leq$ SOG $< 23\,\text{kn}$ ($\hat{p}_j = 0.907$)

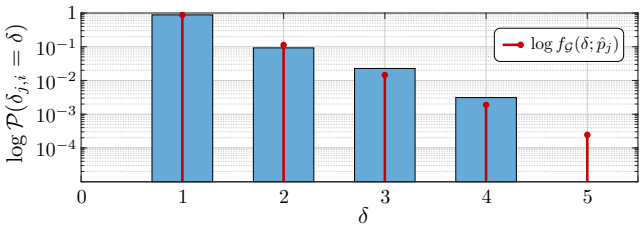

(c) Vessel $j = 5$ at SOG $\geq 23\,\text{kn}$ ($\hat{p}_j = 0.871$)

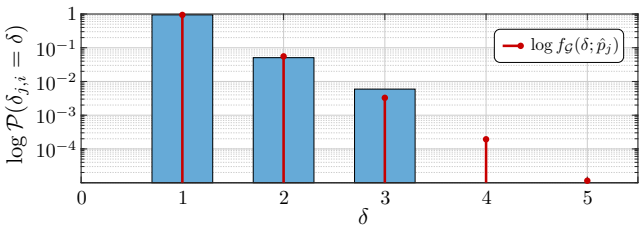

(d) Vessel $j = 6$ at SOG $\geq 23\,\text{kn}$ ($\hat{p}_j = 0.941$)

Fig. 4. Comparison between probability distribution of normalized data and the geometric pmf with parameter $\hat{p}_j$ for different vessels at different SOG intervals.

Our assumption is that the discrete intervals $\delta_{j,i}$, $i = 1, \ldots, N_j$, are i.i.d. random variables geometrically distributed with parameter p_j, i.e., $\delta_{j,i} \sim \mathcal{G}(p_j)$, whose probability mass function (pmf) we denote with $f_{\mathcal{G}}(\delta_{j,i}; p_j)$. The parameter p_j is unknown and can be estimated through the following maximum likelihood (ML) estimator [23, Sec. 9.3]:

$$\hat{p}_j = \frac{N_j}{\sum_{i=1}^{N_j} \delta_{j,i}}.$$

A qualitative way to measure the goodness of fit is given by the root-mean squared error (RMSE). It can be evaluated using the histogram of $\delta_{j,i}$ already reported in Fig. 3c but normalized as a probability distribution, and the pmf $f_{\mathcal{G}}(\delta_{j,i}; \hat{p}_j)$. This comparison is done, for validation purposes, for all the vessels included in our dataset; a small selection of outcomes is reported in Tab. I. Comparisons between empirical and theoretical pmfs for some of these vessels that navigate at

different SOGs are illustrated in Fig. 4. The plots show some differences on the tails; while some discrepancies between the model and the real data are expected,[3] note that these are exacerbated by the use of logarithmic scaling. Indeed, the model is able to capture most of the probability mass, as demonstrated by the experimental results that will be presented in the next section.

V. Empirical Performance Analysis

The geometric distribution appears to model correctly the process under the assumptions that we have described. In this section, we validate the statistical hypothesis procedure by constructing the empirical P_{FA} curve, and subsequently comparing it with the theoretical curve (11).

In Sec. III-A, we have introduced the probability p^*, representing the nominal message reception rate under normal circumstances, and we have mentioned that it can be estimated from the vessels close to that under test. In order to evaluate the empirical P_{FA}, we estimate p^* from the entire real-world AIS dataset shown in Fig. 1 using an ML estimator and assuming $p_j = p^*$, $\forall j$; formally,

$$\hat{p}^* = \frac{\sum_{j=1}^{J} N_j}{\sum_{j=1}^{J} \sum_{i=1}^{N_j} \delta_{j,i}}.$$

It is worth mentioning that the fundamental assumption that all message reception rates from other vessels than the one under test are equal or, in other words, that discrete intervals from vessels that are close in space and time follow the same geometric distribution (cf. Sec. II), is not easily verified in a real-world scenario. This aspect will be discussed in the following, along with possible mitigation methods.

The construction of the empirical P_{FA} requires a sufficient number of trials to evaluate the realizations of the random variable S_K. To this aim, we employ an approach articulated in the following steps.

1) Build a vector $\boldsymbol{X}_j = [X_{j,1}, \ldots, X_{j,K_j}]^{\mathrm{T}}$ of message reception successes (or failures) for each vessel $j = 1, \ldots, J$, starting from the discrete intervals $\delta_{j,i}$ as in (3).
2) Concatenate the vectors $\boldsymbol{X}_j$ for all vessels to obtain the vector $\boldsymbol{\Xi} = [\boldsymbol{X}_1^{\mathrm{T}}, \ldots, \boldsymbol{X}_J^{\mathrm{T}}]^{\mathrm{T}}$ of length $L = \sum_{j=1}^{J} K_j$.
3) Split $\boldsymbol{\Xi}$ into non-overlapping windows of $K < L$ samples. This produces $\lfloor L/K \rfloor$ vectors $\boldsymbol{\xi}^{(w)} = [\xi_1^{(w)}, \ldots, \xi_K^{(w)}]^{\mathrm{T}}$, with $w = 1, \ldots, \lfloor L/K \rfloor$.
4) Compute the decision statistic (8): $S_K^{(w)} = \sum_{i=1}^{K} \xi_i^{(w)}$.
5) Evaluate the empirical P_{FA} by counting how many times, out of $\lfloor L/K \rfloor$ trials, the value of $S_K^{(w)}$ is below γ, for $\gamma \in [0, K]$.

Since the total number of data — and hence the number L — is fixed, in order to obtain a reliable empirical P_{FA} the value K has to be selected as a suitable trade-off between the number $\lfloor L/K \rfloor$ of realizations of $S_K^{(w)}$ and the number K of samples within each window.

[3]One would need an extremely large number of samples to accurately estimate the underlying distribution, especially on the tails.

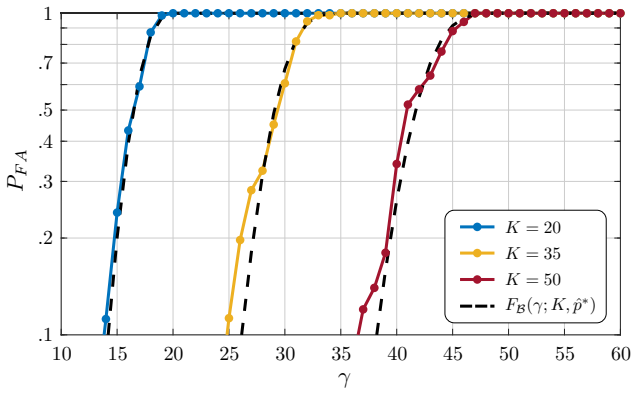

Fig. 5. Comparison between the empirical P_{FA} (colored solid lines) and the theoretical P_{FA} (black dashed lines) as function of the threshold γ, for different values of K. Only vessels with $0.79 \leq \hat{p}_j \leq 0.87$ are considered.

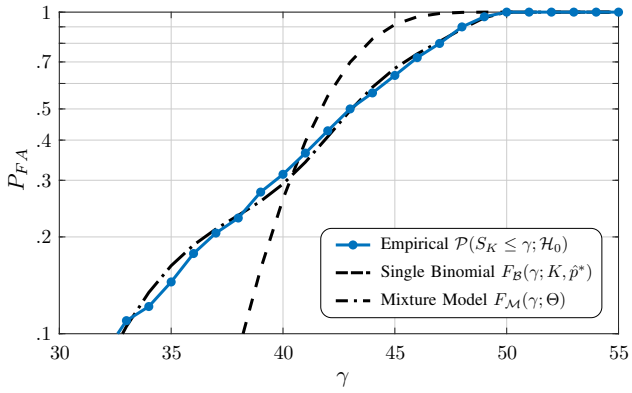

Fig. 6. Comparison between the empirical P_{FA} (blue solid line), the single binomial CDF (black dashed line), and the BMM CDF of 3 components whose parameters are selected to fit the empirical P_{FA} (black dash-dotted line), obtained using the entire real-world dataset.

With reference to Tab. I, we note that the estimated message reception rate can vary significantly from vessel to vessel, with $\hat{p}_j$ spanning between 0.73 and 0.95. Therefore, we focus a first analysis on a few vessels in the dataset whose $\hat{p}_j$ is close to the parameter $\hat{p}^* = 0.83$ estimated from the entire dataset. In other words, we identify the vessels exhibiting the most similar message reception rates to the nominal value, and we apply the procedure just described to compute the empirical P_{FA} curves. Figure 5 demonstrates a very good fit between the theoretical binomial model and the empirical curves; we note that for high values of K (e.g., $K = 50$), the accuracy of the fit decreases due to the lower number of binomial realizations.

A second analysis is devoted to the dataset in its entirety. In this case, strong tails appear in the empirical P_{FA}, which evidently deviates from the single binomial distribution, as shown in Fig. 6. This is not unexpected since, as previously said, the values of $\hat{p}_j$ can vary from vessel to vessel. To deal with this non-ideal situation and come up with a more powerful representation, we can resort to a binomial mixture model (BMM). The downside of using a mixture model is that the choice of the number of components M and their

parameters is not trivial; a rigorous approach would require the use, e.g., of model order selection criteria (see, e.g., [24]). To seek for the best parametric model is beyond the scope of this article, and it may be the subject of future works. The CDF of a mixture of M binomials can be expressed as

$$F_{\mathcal{M}}(\gamma;\Theta)=\sum_{i=1}^{M}\omega_i\cdot F_{\mathcal{B}}(\gamma;n_i,p_i),$$

where $\Theta=[\omega_1,n_1,p_1,\ldots,\omega_M,n_M,p_M]$ is the set of parameters of the BMM. In our case, the components of the BMM all have the same number of independent trials, i.e., $n_i=K,\forall i=1,\ldots,M$. Figure 6 depicts the BMM CDF of three components for $K=50$ and with the six remaining parameters selected to best fit the empirical P_{FA}.

In summary, the experiments performed on real-world data suggest that Dark-VADER is suitable to detect anomalous message delay events in AIS data. When the variability of vessel message reception rates under $\mathcal{H}_0$ is limited, the theoretical model proposed in Sec. II is in agreement with the empirical distribution of message interarrival times (or, equivalently, the sequence of reception successes/failures). When the variability of vessel message reception rates under $\mathcal{H}_0$ is high, the binomial model is not able to follow up with the empirical distribution. Still, a straightforward generalization leads to a more sophisticated model (the BMM), which is able to perform very well also in presence of high variability in message reception rates.

VI. Conclusion

In this paper, we approached the problem of detecting anomalous message delay events in AIS transmissions. To this aim, we proposed the Dark-VADER (dark vessel AIS delay event recognition) algorithm, which models the AIS message reception process as a Bernoulli process and attempts at detecting AIS switch-off events by means of a statistical hypothesis test. The counterpart of the Bernoulli process model is that the time between message arrivals follows a geometric distribution, which was also validated against real-world AIS data, demonstrating that the message interarrival times actually exhibit the postulated decay.

Future advances of Dark-VADER may develop along different research lines. One is the deeper investigation and refinement of the binomial mixture model (BMM); another is the validation of Dark-VADER in more complex real-world scenarios including maneuvering vessels; finally, a third possible line of research is the extension of the algorithm to the dynamic case where the relevant parameters (notably vessel message reception rates) vary in time and space.

References

[1] R. O. Lane, D. A. Nevell, S. D. Hayward, and T. W. Beaney, "Maritime anomaly detection and threat assessment," in *Proc. FUSION-10*, Edinburgh, UK, Jul. 2010, pp. 1–8.

[2] G. Soldi, D. Gaglione, S. Raponi, N. Forti, E. d'Afflisio, P. Kowalski, L. M. Millefiori, D. Zissis, P. Braca, P. Willett, A. Maguer, S. Carniel, G. Sembenini, and C. Warner, "Monitoring of critical undersea infrastructures: The Nord Stream and other recent case studies," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 38, no. 10, pp. 4–24, Oct. 2023.

[3] N. Forti, E. d'Afflisio, P. Braca, L. M. Millefiori, S. Carniel, and P. Willett, "Next-gen intelligent situational awareness systems for maritime surveillance and autonomous navigation [Point of View]," *Proc. IEEE*, vol. 110, no. 10, pp. 1532–1537, Oct. 2022.

[4] *Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile frequency band*, ITU Radiocommunication Sector (ITU-R) Recommendation M.1371-5, Feb. 2014.

[5] M. Svanberg, V. Santén, A. Hörteborn, H. Holm, and C. Finnsgård, "AIS in maritime research," *Marine Policy*, vol. 106, p. 103520, Aug. 2019.

[6] E. Tu, G. Zhang, L. Rachmawati, E. Rajabally, and G.-B. Huang, "Exploiting AIS data for intelligent maritime navigation: A comprehensive survey from data to methodology," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 5, pp. 1559–1582, May 2018.

[7] International Maritime Organization, "International Convention for the Safety of Life at Sea (SOLAS)," Dec. 2000. [Online]. Available: https://wwwcdn.imo.org/localresources/en/KnowledgeCentre/IndexofIMOResolutions/MSCResolutions/MSC.99(73).pdf

[8] "Commission Directive 2011/15/EU," Feb. 2011. [Online]. Available: https://eur-lex.europa.eu/eli/dir/2011/15/oj

[9] T. Stach, Y. Kinkel, M. Constapel, and H.-C. Burmeister, "Maritime anomaly detection for vessel traffic services: A survey," *J. Mar. Sci. Eng.*, vol. 11, no. 6, p. 1174, Jun. 2023.

[10] U. R. Sumaila, J. Alder, and H. Keith, "Global scope and economics of illegal fishing," *Marine Policy*, vol. 30, no. 6, pp. 696–703, Nov. 2006.

[11] E. d'Afflisio, P. Braca, L. M. Millefiori, and P. Willett, "Detecting anomalous deviations from standard maritime routes using the Ornstein–Uhlenbeck process," *IEEE Trans. Signal Process.*, vol. 66, no. 24, pp. 6474–6487, Dec. 2018.

[12] L. M. Millefiori, P. Braca, K. Bryan, and P. Willett, "Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 52, no. 5, pp. 2313–2330, Oct. 2016.

[13] A. Aubry, P. Braca, E. d'Afflisio, A. De Maio, L. M. Millefiori, and P. Willett, "Optimal opponent stealth trajectory planning based on an efficient optimization technique," *IEEE Trans. Signal Process.*, vol. 69, pp. 270–283, 2021.

[14] F. Mazzarella, V. F. Arguedas, and M. Vespe, "Knowledge-based vessel position prediction using historical AIS data," in *Proc. SDF-15*, Bonn, Germany, Oct. 2015, pp. 1–6.

[15] D. Nguyen, R. Vadaine, G. Hajduch, R. Garello, and R. Fablet, "GeoTrackNet—A maritime anomaly detector using probabilistic neural network representation of AIS tracks and a contrario detection," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 6, pp. 5655–5667, Jun. 2022.

[16] M. Vespe, I. Visentini, K. Bryan, and P. Braca, "Unsupervised learning of maritime traffic patterns for anomaly detection," in *Proc. DF&TT-12*, London, UK, May 2012, pp. 1–5.

[17] M. Guerriero, S. Coraluppi, and C. Carthel, "Analysis of AIS intermittency and vessel characterization using a hidden Markov model," NATO-STO Centre for Maritime Research and Experimentation (CMRE), Technical Report NURC-FR-2010-002, Jan. 2010.

[18] F. Mazzarella, M. Vespe, A. Alessandrini, D. Tarchi, G. Aulicino, and A. Vollero, "A novel anomaly detection approach to identify intentional AIS on-off switching," *Expert Syst. Appl.*, vol. 78, pp. 110–123, Jul. 2017.

[19] E. d'Afflisio, P. Braca, and P. Willett, "Malicious AIS spoofing and abnormal stealth deviations: A comprehensive statistical framework for maritime anomaly detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 57, no. 4, pp. 2093–2108, Aug. 2021.

[20] S. K. Singh and F. Heymann, "Machine learning-assisted anomaly detection in maritime navigation using AIS data," in *Proc. PLANS-20*, Portland, OR, USA, Apr. 2020.

[21] *Revised Guidelines for the Onboard Operational Use of Shipborne Automatic Identification Systems (AIS)*, International Maritime Organization Resolution A.1106(29), Dec. 2015.

[22] S. M. Kay, *Fundamentals of Statistical Signal Processing, Detection Theory*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1998.

[23] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*. New York, NY, USA: Springer Science & Business Media, 2004.

[24] P. Stoica and Y. Selen, "Model-order selection: A review of information criterion rules," *IEEE Signal Process. Mag.*, vol. 21, no. 4, pp. 36–47, Jul. 2004.