A Spatio-temporal Track Association Algorithm Based on Marine Vessel Automatic Identification System Data

Imtiaz Ahmed¹, Mikyoung Jun² and Yu Ding³

Abstract—Tracking multiple moving objects in real-time in a dynamic threat environment is an important element in national security and surveillance system. It helps pinpoint and distinguish potential candidates posing threats from other normal objects and monitor the anomalous trajectories until intervention. To locate the anomalous pattern of movements, one needs to have an accurate data association algorithm that can associate the sequential observations of locations and motion with the underlying moving objects, and therefore, build the trajectories of the objects as the objects are moving. In this work, we develop a spatiotemporal approach for tracking maritime vessels as the vessel's location and motion observations are collected by an Automatic Identification System. The proposed approach is developed as an effort to address a data association challenge in which the number of vessels as well as the vessel identification are purposely withheld and time gaps are created in the datasets to mimic the real-life operational complexities under a threat environment. Three training datasets and five test sets are provided in the challenge and a set of quantitative performance metrics is devised by the data challenge organizer for evaluating and comparing resulting methods developed by participants. When our proposed track association algorithm is applied to the five test sets, the algorithm scores a very competitive performance.

Index Terms—AIS, online clustering, threat detection, track association, trajectory tracking.

I. Introduction

N modern days we are blessed with technologically advanced, sophisticated position tracking systems that can transmit the positions and movement of multiple objects in real time. However, we still lack effective spatio-temporal algorithms that can process this large amount of information gathered over time, helping us build, group and predict trajectories, and in some cases, associate tracking points to their true tracks in the absence of object identification information. The idea of associating or assigning unlabeled moving objects to their true tracks is known as *track association*, extremely important for identifying threats in the form of anomalies. Tracking moving objects using both space and time information to detect anomalous trajectory patterns has far reaching safety implications for maritime security. Towards this end, the National Geospatial-intelligence Agency (NGA),

in collaboration with the National Science Foundation (NSF)'s Algorithms for Threat Detection (ATD) program, recently launched data association challenges [1], and design a set of challenge problems using the maritime vessel tracking data, collected by the vessels' Automatic Identification System (AIS). The research reported in this paper is an algorithm developed to address the ATD data association challenge problems. We want to stress that the focus of this work is on associating the AIS measurements/detections to true vessels while the measurements have already been collected by the AIS. We will explain the datasets in more detail in Section II.

AIS is an automated maritime vessel locating, tracking, and monitoring system. Each ship is equipped with a transponder and uses a common very-high-frequency radio channel to communicate with other vessels, the shore-based stations, and some global positioning satellites. AIS was developed under the guidance of the International Maritime Organization (IMO) and is required on most ocean-going commercial ships. According to IMO 2002 convention, international voyaging ships, with 300 or more gross tonnage, and all passenger ships regardless of size require installation of AIS. The primary purpose of AIS used to be helping avoid unwanted collisions among vessels and thereby enhance maritime safety. Nowadays, however, AIS is increasingly used as a monitoring system to track vessel movements at the open sea or near a port [2], because AIS can locate another vessel more effectively under reduced or zero visibility than the radar detection system, and thus enable watch guards to prepare in advance [3]. For this reason, AIS also helps with maritime traffic management and the maintaining of coastal security.

AIS data transmitted by each vessel are in the form of time-sequenced nodes. Each node represents a row or a data point in the AIS data file and contains the timestamps (timestamps marking the signal receiving time of the AIS receiver), coordinates (latitude and longitude), speed (in knots), and direction of a vessel (in degrees). Time sequenced nodes enable receiving groups of messages from an input source, and preserve the order in which the messages in each group arrived. AIS data also include the vessel's maritime mobile service identity (MMSI) number, unique for each ship. Despite the technological advancement, in the reality of AIS data transmission, collection, and retrieval, the MMSI number could be missing or messed up. When MMSI is absent, it presents a particular challenge for track association. Simply put, the question is that when the MMSI number is removed or missing, could one still associate these time-sequenced nodes

¹Imtiaz Ahmed with the Department of Industrial & Management Systems Engineering, West Virginia University, Morgantown, WV. Email: imtiaz.ahmed@mail.wvu.edu

²Mikyoung Jun with the Department of Mathematics, University of Houston, Houston, TX. Email: mjun@central.uh.edu

³Yu Ding with the Department of Industrial & Systems Engineering, Texas A&M University, College Station, TX. Email: yuding@tamu.edu

correctly in an AIS dataset to recreate the trajectories of the vessels?

Apart from the absence of MMSI, there exist other issues further complicating track association. One such issue is the presence of time gaps in the vessel reporting system. A vessel can randomly stop sending signals for various reasons, ranging from sudden equipment failure to deliberate hiding of the current trajectory. It would be much difficult to associate the nodes to the right track after a long absence of its signal, because during this period, a ship can alter its direction completely, increase the speed dramatically, or even stop altogether. Another difficulty in track association arises when the vessels are near a port. In a port, lots of vessels are parked nearby, making them hard to be differentiated. They also maneuver frequently for parking or for making space for other vessels. Therefore, tracking individual vessels near ports needs special care.

The academic contributions of this work can be summarized as follows:

- ☐ We propose a spatio-temporal track association algorithm that can associate vessel measurements collected from the AIS to their true tracks in real-time.
- □ We propose a novel dissimilarity score metric that compares the predicted location of existing vessels with the current AIS measurement in hand and decides between opening a new track versus assigning the node to an existing track in a complete online fashion. We develop our location prediction algorithm based on Vincentry's direct geodesic framework [4].
- ☐ To address the sensitivity of the online, real-time association to some complex tracking phenomena (e.g., abrupt direction change, parking maneuver near the parking port, absence of AIS signal for a significant period of time etc.), we propose a post-hoc track merging step to prevent the creation of unnecessary new tracks. These thresholds are empirically learned and tuned using the training datasets provided by the challenge organizers.
- ☐ We also propose a spatially varying track generation and merging policy. It divides the underlying surveillance area into three spatial regions and accordingly approves new track generation or suggests possible merging.

Our proposed approach has the following strengths. First, it is online in nature, meaning that it associates each time stamped node with a vessel on the fly. Second, it makes use of the spatial information to increase its effectiveness for tracking vessels in the open sea and near a port. Third, it can track vessels when there are time gaps in the AIS signals. Fourth, it embodies a provision for offline correction of some difficult-to-separate tracks. Our proposed method is applied to the five data sets provided by the data association challenge organizers. The performance metrics, as also designed by the challenge organizers, are calculated for the proposed method. There were more than ten teams participating in the ATD data challenge. According to the information made available to us, our algorithm performs the second best as compared with those teams. But we are not authorized to disclose the full list of performance statistics of other teams. In Section V, we therefore present the performance of our proposed method together with that of a sample baseline algorithm provided by the challenger organizers. We also present the tracking and association performance of two popular multi-object tracking approaches and compare their performances with our proposed approach.

The rest of the paper unfolds as follows. Section II describes the data format, variables and characteristics of the training and test data sets. The performance metrics designed will be explained later in Section V-A. Section III highlights some of the existing research areas relevant to the track association problem based on the AIS or trajectory data. In Section IV, we present the main idea of our track association algorithm. Section V explains the performance evaluation criteria and presents a comparison study of different methods under the same challenge. Finally, we summarise the paper in Section VI.

II. DATA DESCRIPTION

In this section, we describe the AIS data sets [1] used in the ATD data association challenge. The data comes from the historical database of the Nationwide Automatic Identification System supplied by the United States Coast Guard. The data covers the area around Norfolk, Virginia. Six variables are provided for track association, namely, VID (vessel ID), timestamp (in hh:mm:ss format, where hh represent hours, mm represent minutes, and ss represent seconds), latitude (in degrees), longitude (in degrees), speed (in tenths of knots), and course of direction (in tenths of degrees). As speed and direction are expressed in tenths of their respective units, one needs to divide the raw data by 10 to get the actual value. The VID is referred to as the MMSI number in the preceding section.

When the AIS information is received, a monitoring officer can pinpoint it to a single, unique vessel using the VID, if the VID is present and valid. Fig. 1(a) and 1(b) show an illustration under such circumstance. In Fig. 1(a), each letter (A, B, C or D) represents a unique vessel and the number (1, 2, 3, etc.) that identify the time stamped nodes. Each of these nodes has the associated speed, moving direction, and location in terms of longitude and latitude. Fig. 1(b) presents another view, where each of the time stamped nodes is classified into a track associated with a specific vessel. The organization in Fig. 1(b) clearly shows the number of unique vessels and how the time stamped information is arranged along each vessel's moving trajectory.

The AIS data received by ships and coastal stations are then transmitted to regional or national data centers. When multiple receivers are connected into the communication and command network, a number of issues could happen, including data intermittency, data redundancy received by multiple receivers, errors in timestamps assigned by varying receivers, and vessels erroneously sharing the message identifier [5]. When the VID information (letter and color code) is missing from the time stamped nodes (see Fig. 1(c)), the nodes are then numbered according to the time they are generated. The numbering is sequential in time, with "1" meaning the earliest reported node

and "14" meaning the latest reported node in this example. For example, both A1 and A2 represent the trajectory points of vessel A but A1 appears or reported earlier than A2. Then the desire is to develop an algorithm that can analyze the node information without VID, as in Fig. 1(c), and hopefully group the nodes as in Fig. 1(d), which, if done so, can reveal the same vessel and trajectory information as Fig. 1(b) does.

All data files are provided in the comma-separated-values (CSV) format. There are three training datasets, for which the VID associated with each timestamp is given, so that the participants can evaluate the accuracy of their algorithms and then fine tune them. To test the performance of the algorithms generated by the participants, five test datasets are provided by the challenge organizers. The only difference between the training and test datasets is the absence of VID in the test sets. The VID information is intentionally removed so that the competing algorithms can be tested for their association accuracy.

For the purpose of evaluation, competition participants are asked to generate a separate column/variable in the reporting CSV file, known as TRACK_ID which provides the numbers generated sequentially for each timestamp but unique to a vessel. For example, if we assign the TRACK_ID '1' to the first record/timestamp and it happens to represent vessel '100008', then each timestamp of the same vessel but appearing later should also be given the TRACK_ID of '1', if an algorithm works successfully. For each test set, participants are supposed to provide a CSV file with this one extra column of TRACK_ID.

We summarize the characteristics of the training and test datasets in Table I. The training set is referred to as dataset 0 and is further divided into three datasets, i.e., 0-1, 0-2, 0-3, each of which consists of four hours worth of data collected from three separate days. In the training sets, VIDs are known and available to the participants so that they can tune their algorithm parameters to set at optimal values.

Test sets, i.e., datasets 1 to 5, are organized in the increasing order of difficulty. Although the AIS data collection area is somewhat fixed for all five test sets, the track association problem becomes more complicated with the increasing number of vessels and more observations over the five test sets. To add another layer of intricacy, data gap is intentionally incorporated in test sets 4 and 5 by the challenge organizers. This is to mimic the missing data situation often present in real-life datasets. The presence of data gap implies that a selected chunk of the data (e.g., a 20-minute period or several smaller amounts of time for one or more tracks) are removed. All tracks present immediately prior to a gap are still present immediately after the gap, i.e., the track continuity is maintained across each data gap. For the test sets provided, tracks are expected to keep the same IDs on either side of the gap, i.e., no swaps, breaks or new tracks immediately before and after the gaps. Although there are no track terminations or new initiations immediately across a given gap, such termination and initiation are possible at other times.

The gap characteristics are a bit different between test sets 4 and 5. In test set 4, gaps are created for ALL tracks at two fixed

time stamps (not disclosed to the participants), which means that after these particular timestamps, for a certain amount of time no data has been reported from any of the existing vessels. On the contrary, in test set 5, gaps are created only for a subset of the tracks (we do not know which tracks have time gaps). Unlike in set 4, set 5's gaps are not at the same time. Assumingly, it makes the test set 5 more complex than the test set 4. Furthermore, for the test sets, the participants are given a rough estimate of the number of vessels present, not the exact number, and sometimes, only the lower bound.

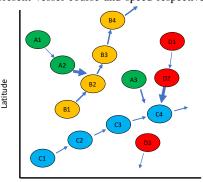
We want to stress that the datasets are given to us by the Data Challenge Organizer (NGA and NSF). We do not include or incorporate any feature in the dataset by ourselves. We just applied our algorithm to the test sets and then compute the performance metrics which are designed by the challenge oganizers as well. We can guess that they add different features (e.g., time gap, more vessels, more observations) to the test data sets to make the challenge problems more complex. In other words, the Data Challenge specifies strictly what are the inputs and what are the outputs (performance metrics), and all they want is a newer and more capable algorithm in between.

The target of the AIS data association challenge algorithm is to develop new spatio-temporal association algorithms that can track multiple vessels in a difficult to track scenario and associate the nodes correctly to their respective tracks. Some of the challenges discussed above are purposely introduced so that the existing approaches become much less effective in the association tasks. To provide a quick summary, we further list the following characteristics, which are considered as the bottleneck of the data association challenge:

- Each data point is one real object, not a clutter.
- Vessels can pass each other at a very close distance as some of the data corresponds to their parking maneuver in the port.
- The vessel's tracks can cross each other on multiple occasions making it harder to associate the nodes to their corresponding tracks correctly.
- The number of vessels is unknown. Multiple vessels information can be produced at a single time point. The frequency of the incoming information is not fixed and fluctuates through out the information collection period.
- Time gaps are incorporated in the AIS data which means that some of the timestamps are purposely removed from the AIS data. A vessel's trajectory can undergo significant changes during this time gap period thus making the job even harder for the association algorithms.

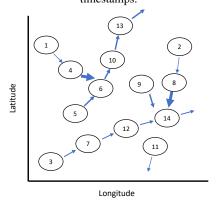
III. LITERATURE REVIEW

There has been a growing interest over the years to collect and analyze the AIS and marine surveillance data to extract important knowledge from the trajectory pattern. These research works can be categorized into a few major schools of thought, namely route estimation or prediction, trajectory clustering, anomalous pattern identification and multi-object, multi-sensor tracking. (a) Graphical depiction of the AIS data. Each letter and color code represent an unique vessel. Direction and width of the arrows represent vessel course and speed respectively.

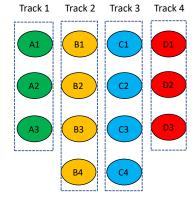


Longitude

(c) The AIS data when vessel ID is removed. The numbers represent the objects IDs generated sequentially with respect to their timestamps.



(b) Tracks consisting of nodes grouped by vessels.



(d) Track association in absence of vessel ID.

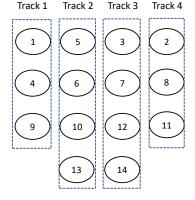


Fig. 1: Illustration of the track association problem.

TABLE I: Characteristics of the training and test sets

Problem set	VID included	AIS data	Number of	Number of	Data gap
		duration	vessels	observations	
0-1 (training)	Yes	4 hours	20	13,714	No
0-2 (training)	Yes	4 hours	26	15,707	No
0-3 (training)	Yes	4 hours	23	12,940	No
1	No	4 hours	8-10	3,323	No
2	No	4 hours	8-10	8,056	No
3	No	4 hours	>10	16,810	No
4	No	4 hours	>10	20,538	Yes
5	No	4 hours	>10	24,513	Yes

Route Prediction

Route prediction refers to the prediction of the trajectory pattern of a vessel over time. It includes the estimation of future position and other motion characteristics [5]. It is one of the widely explored research areas involving the AIS data. Route estimation can be further classified into three classes based on their estimation approaches: the physics-based models (motion models) [6, 7, 8], the learning-based models [9, 10, 11, 12], and the hybrid models [13, 14, 15, 16]. It is important to keep in mind that the movement of maritime surface vessels have some unique characteristics, differentiating them from land, air, or even underwater vehicles. For instance, a water surface vessel cannot abruptly stop or change direction [17]. It spends more time and covers more space during a motion-changing process. Also, a surface vessel

moves in a two-dimensional horizontal plane more like a land vessel but unlike air or underwater vessels, which can also move vertically. For the same reason, the tracking task for marine vessels is also different from typical image based object tracking [18] or video tracking [19]. Luckily, some of these unique characteristics do help predict the route of maritime surface vessels.

Trajectory Clustering

Trajectory clustering is to group similar trajectories into clusters. For any clustering model, the single most important factor is to decide on which similarity measure to be used for the purpose of clustering [20]. A trajectory is composed of a series of multi-dimensional spatio-temporal points with irregular sampling interval. Traditional similarity measures

somewhat fall short of handling this arbitrary discretization [21]. As a result, devising various similarity measures, better fit to the vessel trajectory problems, have been a main focus of research effort over the years [22, 23, 24]. Some popular choices are the principal component analysis (PCA) plus Euclidean distance, Hausdorff distance, Fréchet distance, longest common sub-sequence distance (LCSS), and dynamic time warping (DTW). Different algorithms have been tried on trajectory data, including K-means [25], balances iterative and clustering using hierarchies (BIRCH) [26], density-based spatial clustering of applications with noise (DBSCAN) [27], and statistical information grid (STING) [28]. These trajectory clustering methods can be classified into five groups: spatialbased clustering [29, 30, 31, 32], time-dependent clustering [33, 34], partition and group-based clustering [35, 36], uncertain trajectory clustering [37], and semantic trajectory clustering [38].

Anomalous Pattern Identification

Identification of anomalous patterns from the vessel trajectory data helps locate suspicious activities in the sea [39]. To detect anomalous patterns, the first step is to establish a normalcy baseline, so that a deviation from the baseline can be signaled for further investigation [40]. The current algorithms for anomalous trajectory pattern detection rely mostly on finding the most representative trajectory from a pool of trajectories, treating it as the normalcy baseline, and then comparing each trajectory with the baseline using a similarity measure [41, 42]. Anomalies can be tracked to a specific region of the trajectory and features responsible for them. Understandably, similar to the trajectory clustering approaches described above, a suitable dissimilarity measure plays a crucial role in comparing trajectories and signaling potential deviation.

Multi-object Trackers (MOT)

For maritime security and surveillance, multi-objects trackers (MOT) have also been useful which fuses information from multiple sensors such as radar and sonar rather than directly utilizing the AIS data. MOTs are capable of estimating the future position of moving objects using the past positions of these objects and doing subsequent association. Several tracking algorithms are proposed along this line. Global nearest neighbor (GNN) [43] is the simplest algorithm that assigns measurement to track by minimizing a cost assignment matrix. Joint Probabilistic Data Association (JPDA) [44] is another widely used tracking algorithm where the tracker uses a soft assignment (track association probability) and therefore multiple objects can contribute to the update of each track. Both GNN and JPDA maintain a single hypothesis about the tracked objects. Multiple Hypotheses Tracking (MHT) [45] is another popular approach for tracking and association which generates a tree of hypotheses for each object. The likelihood of each track is calculated and the most likely combination of tracks is selected. MOT algorithms generally rely on Kalman filter [46] to estimate the vessel's dynamic states from positional and directional measurements.

However, MOTs, though seemed a good candidate for vessel tracking and association, are not ideal for handling our AIS challenge datasets. The first issue is that we have cleaned and clutter free AIS object (vessel) information in hand instead of noisy information from sensor(s) that needs to be processed and fused by the MOTs. Second, we can utilize the spatial information of the underlying surveillance area which is different from these conventional MOTs. Third, we believe that the tracking filters used in MOT algorithms are not capable of handling the complex tracking phenomena present in these challenge datasets (e.g., non-linear trajectory, abrupt direction or velocity change between two updates, parking maneuver near the parking port, absence of AIS signal for a significant period of time etc.). When we apply GNN and JPDA to our challenge datasets, they produce poor track association performances (detail in Section V-C).

IV. THE PROPOSED TRACK ASSOCIATION ALGORITHM

In this section, we propose our approach for the track association problem. Our approach consists of two separate stages: (1) Online track association and (2) Post-hoc merging. We want to first elaborate the design of our approach, especially the thought process behind the two stages. At the high level, we, in both stages utilize the unique movement features of the sea vessels (more or less steady course, no abrupt changes in speed and direction) to create detection conditions, and then use empirically learned thresholds to check these conditions and guide track association.

The first stage works in an online manner. It associates each node with a track as it appears along the process. For each node, our online track association algorithm either chooses to open a new track or assigns the node to an existing track. It is built upon a location prediction framework which enables us to estimate the next node location of an existing track and thereby helps anticipate the current node position for deciding its track association. This prediction fits perfectly into the famous Vincenty's direct geodesic problem framework [4]. For that, we make use of a recent solution [47] of the geodesic problem to estimate the node location. For the association task, we develop a novel dissimilarity metric that helps us to find the right association between a track and the underlying AIS measurement. The solution works quite well when the AIS data is collected at a higher frequency, i.e., the distance traveled by a vessel before the receipt of its next AIS measurement is small, so having this competent solution adds the first competitive edge to our online association approach.

Furthermore, to make our approach robust to real-life operational complexity, we devise a varying threshold policy which takes into consideration the sensitivity of the prediction formula conditioned on the distance a vessel has recently traveled. We find that this conditional threshold policy injects the second competitive edge to our online association.

Despite the much enhanced capability of our online association approach, it is apparent to us that the online algorithm could still make a number of mistakes, especially in terms of creating unnecessary new tracks, due to its myopic nature. This observation motivates us to design a second stage, after

completing the online stage, for deciding whether some of the tracks should be merged. Particular attention in the second stage is given to the spatially varying pattern of new track generation. We adopt policies, learned from the training data, to treat special cases, e.g., significant time gap between two nodes from the same track, and vessels getting in or out of parking locations. The inclusion of this location-sensitive posthoc merging in the second stage is the third reason behind the success of our overall track association algorithm.

A. Online Track Association

Let us represent the set of true tracks as I, the set of associated tracks as J, and the set of incoming nodes/timestamps as K. Understandably, the size of set I and that of set J may not be equal. Each node/timestamp, $k \in K$, comes with the AIS measurements, namely the time of measurement (t_k) , current position (p_k) , which has two components, the latitude, ϕ_k , and the longitude, λ_k), speed (v_k) , and course of direction (θ_k) . Note that we convert all speed values from knots to meters per second (m/s).

The online track association stage starts off by opening an associated track (j=1) and assigning the very first node (k=1) to this track. From this point onward, for each incoming node, k, we either assign it to one of the existing tracks in set J or open a new track, i.e., add a new element to J. We use the notation, k^j , to indicate that node k is on an associated track j. Following the notation, the AIS measurements of this node k on track j are denoted by t_{k^j} , $p_{k^j} := (\phi_{k^j}, \lambda_{k^j})$, v_{k^j} , θ_{k^j} . When both an associated track and a true track appear, we usually use j to denote the associate track and i to denote the true track. Because of this, unless otherwise indicated, k^i is reserved for node k on a true track i.

To carry out this online evaluation, we calculate an individual dissimilarity score, denoted by s_{jk} as in (1), to assess the similarity between an existing associated track j and the current node k in consideration. A low score implies high similarity and therefore indicates a possible association to an existing track. For each existing associated track, $j \in J$, its estimated next node measurements are indexed by n^j , where n implies next. The most recent AIS measurements on track j is one step back from n and thus represented by $n^j - 1$.

Using these notations, we propose the following dissimilarity score:

$$s_{jk} = c_{dist}(p_k, p_{n^j}) + c_{ang}(\theta_k, \theta_{(n^j - 1)}).$$
 (1)

The two terms, c_{dist} and c_{ang} , denote changes in distance and angle, respectively, and they are calculated as

$$c_{dist}(p_k, p_{n^j}) = 2r \times$$

$$\arcsin \sqrt{\sin^2 \left(\frac{\phi_k - \phi_{n^j}}{2}\right) + \cos \phi_k \cos \phi_{n^j} \sin^2 \left(\frac{\lambda_k - \lambda_{n^j}}{2}\right)}$$

and

$$c_{ang}(\theta_k, \theta_{(n^j - 1)}) = \frac{180 - \left| 180 - \left| \theta_k - \theta_{(n^j - 1)} \right| \right|}{\left| t_k - t_{(n^j - 1)} \right|}, \quad (3)$$

where r denotes the Earth's radius. The c_{dist} , in (2), captures the spatial distance between the predicted location and the

position of the current node based on the Haversine distance [48], whereas the c_{ang} , in (3) measures the angular change of direction over time between the associated track's last known direction and the direction of the current node and it is devised so as to account for the circular nature of the angular measurements, i.e., the 0 and 360 degrees are the same. The spatial proximity (c_{dist}) of the current node and the predicted location of existing vessels helps us to decide whether they belong to the same track or not. On the other hand, c_{ang} tracks the required change in course per unit time to reach to the current location for each existing vessel and thereby indicates the feasibility of grouping them under the same track. It is normally assumed that a vessel does not abruptly change its direction. Together they provide a complete picture of the similarity between the current node and existing tracks.

While most of the variables in (2) and (3) are readily available from the AIS measurements, the location coordinates associated with n^j , i.e., $p_{n^j} := (\phi_{n^j}, \lambda_{n^j})$, need to be predicted using its last node's location, $p_{(n^j-1)}$, and bearing, $\theta_{(n^j-1)}$. This is done by solving the Vincenty's direct geodesic problem [4]. We skip the details of the derivation but present a graphical illustration in Fig. 2. The final formula for the two coordinates are:

$$\phi_{n^j} = \arcsin(\sin\phi_{(n^j-1)}\cos\delta_{jk} + \cos\phi_{(n^j-1)}\sin\delta_{jk} \times \cos\theta_{(n^j-1)})$$
(4)

anc

$$\lambda_{n^j} = \lambda_{(n^j - 1)} + \arctan(\sin \theta_{(n^j - 1)} \sin \delta_{jk} \cos \phi_{(n^j - 1)},$$
$$\cos \delta_{jk} - \sin \phi_{(n^j - 1)} \sin \phi_{n^j}). \tag{5}$$

In the above two equations, to measure the angular distance, δ_{jk} , we let

$$\delta_{jk} = \frac{d_{jk}}{r},\tag{6}$$

where d_{jk} represents the distance traveled along the shortest path on an ellipsoid (the geodesic). We further estimate this d_{jk} through

$$d_{jk} = \frac{v_k + v_{(n^j - 1)}}{2} \times |t_k - t_{(n^j - 1)}|. \tag{7}$$

Note that, the use of the average speed in (7) is an approximation, since the vessel's speed during the time traveled may be quite different (e.g., acceleration, deceleration etc.). Without any additional information to inform us of the vessel's behavior between the two measurements, we find that the use of the average speed provides a good approximation when the distance traveled is small. In this paper, we use a R package called geosphere [49] to generate the values of p_{nj} and calculate the Haversine distance (c_{dist}) . The package only needs the last known measurements of each vessel, current AIS measurement and the distance traveled to reach to the current location.

Finally, we calculate an overall dissimilarity score for node k

$$s_k := \min_{j \in J} s_{jk},\tag{8}$$

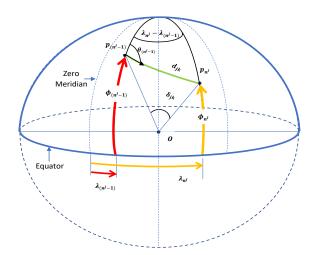


Fig. 2: Illustration of predicting the new node location. The green segment represents the geodesic path followed to reach the new location.

which is the minimum of the individual dissimilarity scores s_{jk} as in (1) over all j's in the set of J. We initially decide to compare s_k against a threshold of β . If $s_k > \beta$, a new track would open, and otherwise the node would be assigned to an existing track, say $z \in J$, which returns the lowest dissimilarity score. That is, $z = \arg\min s_{jk}$ for each k.

We find that using the dissimilarity score, s_k as in (8), there are a couple of complexities that need to be taken are of. The first issue is a result of having s_k as a summation of both the distance change and the direction change. But recall that the maritime surface vessels are not supposed to change its course of direction abruptly, yet s_k , due to its design, is not always sensitive to an angular change alone. For instance, near a port, when vessels are passing each other in close distance with similar speed, then only directional difference can separate these vessels, not the distance and speed. To handle this complexity, we decide to add a threshold, α , only on the directional change, c_{ang} , in addition to checking s_k . When $c_{ang} > \alpha$, then open a new track, regardless of what s_k is; otherwise, chose one of the existing tracks based on s_k .

The second issue is how to choose the threshold for s_k . When s_k is very small or very big, the decision of choosing an existing track (s_k small) or opening a new track (s_k large) is reasonably robust, much as expected. When s_k is in between, we observe that a single threshold on s_k , as we originally envisioned, is not robust enough. Another factor, which is the distance traveled, i.e., d_{zk} as in (7), affects the track association decision. Generally speaking, a large value of d_{zk} usually leads to a relatively less accurate location prediction, and as such, needs a higher s_k value for reaching a better decision.

In light of the above discussion, we create a global minimum value of β , denoted as β_{small} , and a global maximum value, denoted as β_{large} , so that if $s_k \leq \beta_{small}$, we are going to assign node k to track z, i.e., using an existing track, whereas if $s_k > \beta_{large}$, we are going to open a new track. When $\beta_{small} < s_k \leq \beta_{large}$, we devise another layer of decision,

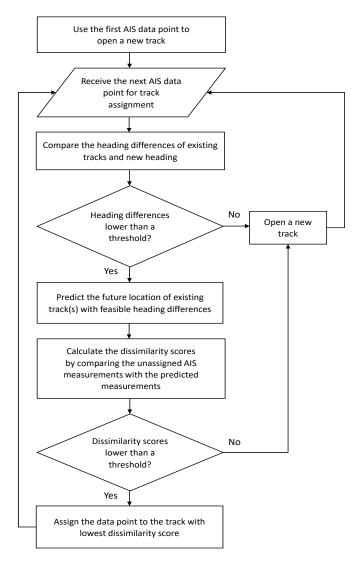


Fig. 3: A schematic flow diagram of the online track association process.

conditioned on d_{zk} . To do that, we introduce another threshold of μ . When $\beta_{small} < s_k \leq \beta_{large}$, we open a new track if $d_{zk} \leq \mu$, otherwise the node would be assigned to the existing track, z.

All of these thresholds, i.e., β_{small} , β_{large} , μ , and α are determined empirically from the training datasets. For better understanding, a schematic flow diagram of the online track association process is summarized in Fig. 3.

B. Post-hoc Merging

The purpose of the post-hoc analysis is to handle the situations that the online clustering does not handle adequately, often due to the need for the online training to keep up with the ongoing process and the limitation of not seeing all the data. Our online clustering method tends to produce more tracks than the number of true tracks, so that our post-hoc analysis does mostly merging than splitting.

To consider the possibility of merging, we create a set of candidate tracks, $A^{(j)}$ for each associated track j, where $A^{(j)} \subset J$ and $j \notin A^{(j)}$. Let us represent the start and end node of an associated track $j \in J$ as s^j and e^j , respectively; again the superscript j here signifies the association of s and e with track j. Given a track $a \in A^{(j)}$, the end node of track a is thus denoted by e^a . For the subsequent analysis, we restrict the tracks in $A^{(j)}$ to be those that stopped before track j, i.e., $t_{s^j} > t_{e^a}$ for $a \in A^{(j)}$.

Our post-hoc merging targets two major problems. The first problem is the wrongful track association due to the presence of time gap. In the presence of such time gap, our location prediction becomes less accurate, and consequently, the online clustering step tends to open a new track whenever a new node appears after the gap, even if it is from an existing track. We consider the following situation as candidates for merging—a candidate track whose last node appeared more than τ seconds apart from the first node of the new track, yet the two nodes are close than a distance threshold γ . The condition is formally stated below:

$$t_{s^j}-t_{e^a} \geq \tau \text{ and } c_{dist}(p_{s^j},p_{e^a}) \leq \gamma, \quad \text{for } a \in A^{(j)}. \quad (9)$$

The second issue to handle is when a vessel changes its direction abruptly, i.e., when $c_{ang}(\theta_k,\theta_{(n^j-1)})$ is large, which leads the online step to open a new track (recall the $c_{ang}>\alpha$ criterion used in the online clustering step). While we institute the angular change condition to open a new track, it could open too many new tracks. But with the angular condition in the online step, the algorithm will miss genuine new tracks. So it is a balance act between over-detection and under-detection. We find a simpler approach, which is to have the online angular condition as it is but do a post-hoc correction. The correction is to merge tracks which have their start and end nodes too close to each other, expressed below:

$$c_{dist}(p_{s^j}, p_{e^a}) \le \eta. \tag{10}$$

Note that whatever issue arises, merging will always take place with a single candidate track, a. If there are multiple choices, we will always select the nearest track.

We observe that the likelihood of a post-hoc merging depends heavily on the nature of the locations where a new track emerges. For the datasets provided to the participants, they cover a finite area nearby the Norfolk port. The vessel locations are truncated, so that there is a clear boundary surrounding the area under monitoring where the AIS data are available; please see Fig. 4. We notice that three types of locations provides important clues for merging likelihood: the boundary locations (in the open sea), the middle locations, and the parking location (which is the boundary locations on land). We devise the following rules of thumb to guide the post-hoc merging based on some common understanding:

 Middle locations: New tracks cannot appear suddenly out of nowhere. When a new track appears here, it is either the result of a time gap or due to an abrupt direction change in the course of an existing track. We must merge such tracks with the nearby track.



Fig. 4: Illustration of boundary, parking and middle locations for track initiation using test set 5. Black line indicates the boundary and the red lines indicates the vessel positions.

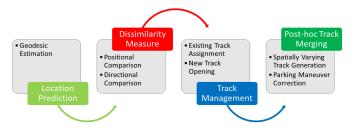


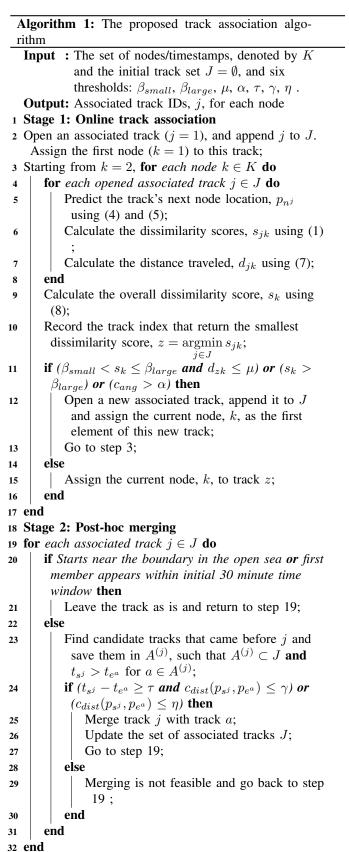
Fig. 5: A block diagram to summarize the proposed framework.

- Boundary locations: New tracks could emerge any time
 at the boundary points. As there is no information discerning where the new track comes from, we cannot rule
 out the possibility of a genuine new track, so such tracks
 will not be evaluated for merging.
- Parking locations: New tracks can initiate from the parking locations but the historical data show that the initial timestamps of such tracks are always within the first 30 minutes. So we will test a new track for merging only when it emerges after 30 minutes.

Combining all the elements, the whole algorithm is summarized in Algorithm 1. We also summarized the key components of the entire association process in Fig. 5.

V. PERFORMANCE EVALUATION

In this section, we evaluate the performance of our proposed track association algorithm on five test datasets. In Section V-A, we describe the performance metrics used for the performance evaluation. In Section V-B, we provide the detailed performance of our approach, together with the sample baseline algorithm provided for the data association challenge. In Section V-C, we evaluate the performance of two traditional MOT approaches with respect to our proposed track association algorithm. Finally, in Section V-D, we evaluate the performance of our proposed approach on an open-source AIS dataset.



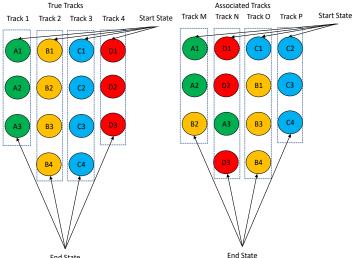


Fig. 6: A sample scenario of true and associated tracks, where the start and end states of each track are marked.

A. Performance Metrics

To assess and quantify the performance of a tracking algorithm, the ATD data challenge organizers provided three performance metrics [1], which are described in the sequel.

1) Counts of erroneous tracks: The purpose of this metric is to measure the quality of track association by counting the number of instances of different errors that could happen during the association. To provide a clear understanding, a simple track association example is presented in Fig. 6. It has four true tracks (Track 1, Track 2, Track 3, Track 4). Each track consists of multiple nodes, where a pair of consecutive nodes form a segment, e.g., A1 and A2 are the two nodes from Track 1 and A1-A2 is a segment. There are 14 nodes in total, which are distributed into the four tracks. There are also four associated tracks (Track M, Track N, Track O, Track P), resulted from a track association algorithm. As this example is generated to depict different errors, the track association is arbitrarily set and understandably made imperfect. The start and end states of each of these tracks are specifically marked as they play a major role in defining different errors.

We use five error categories to check the quality of a track association algorithm: missed track, extra track, merged track, broken track, and swapped track. Other than the aspect of swapped tracks, we use only the start and end states of the associated tracks to measure these errors. They are defined below:

- Missed track: A true track's start state that is not captured by any of the associated track's start states. In the sample example, node B1 corresponds to a missed track.
- Extra track: An associated track's start state that does not match with any of the true track's start states. In the sample example, node C2 corresponds to an extra track.
- Merged track: A true track's end state that is not captured by any of the associated track's end states. In the sample example, node A3 corresponds to a merged track.
- Broken track: An associated track's end state that does

not match with any of the true track's end states. In the sample example, node B2 corresponds to a broken track.

As one only needs the start and end states, these four error categories are also applicable when the number of true and associated tracks are not the same. We do not use the internal nodes in counting these four errors to keep the evaluation process as simple as possible.

To count the number of swapped tracks we have to use the information of track segments. A track swap occurs when a segment on a true track is not on any of the associated tracks. From Fig. 6, ten true track segments can be identified: A1-A2, A2-A3, B1-B2, B2-B3, B3-B4, C1-C2, C2-C3, C3-C4, D1-D2, and D2-D3. Similarly, ten associated track segments can be defined: A1-A2, A2-B2, D1-D2, D2-A3, A3-D3, C1-B1, B1-B3, B3-B4, C2-C3, and C3-C4. We can see that there are five true track segments (A2-A3, B1-B2, B2-B3, C1-C2, and D2-D3) which are not on any of the associated tracks. So, the number of track swaps in this case is five. The lower bound for all these errors is zero, which represents a perfect association. Note that using the five errors, as described above, does not present a clear indication of track matching or association. Track association will be done using the completeness score, to be described in Section V-A3.

2) Continuity score: A continuity score is conceptually the mirror of track swapping. It is defined as the ratio of the distance of all the correct associated track segments to the distance of all the true track segments:

Continuity score
$$= \frac{\sum_{j \in J} \sum_{k^{j}=2}^{n_{(k^{j})}} g(p_{(k^{j})}, p_{(k^{j}-1)})}{\sum_{i \in I} \sum_{k^{i}=2}^{n_{(k^{i})}} c_{dist}(p_{(k^{i})}, p_{(k^{i}-1)})}$$
(11)

with

$$g(p_{(k^{j})}, p_{(k^{j}-1)}) = \begin{cases} c_{dist}(p_{(k^{j})}, p_{(k^{j}-1)}) & \text{if } [k^{j}, (k^{j}-1)] \in T \\ 0 & \text{if } [k^{j}, (k^{j}-1)] \notin T \end{cases}$$
(12)

Here, k^i and k^j denote running indices for the nodes of true track i and associated track j, respectively. Furthermore, $n_{(k^i)}$ and $n_{(k^j)}$ represent numbers of total nodes available under true track i and associated track j, respectively. The term T represents the set of true track segments. We use the Haversine distance measure, c_{dist} , for each true segment as defined in (2). The continuity score is normalized between 0 and 1, where 1 means perfect. Similar to the preceding five error counts, using the continuity score does not lead to track matching or association, either.

3) Completeness score: Completeness of a true track with respect to an associated track is measured by taking the ratio of the number of nodes that are present in both the concerned true and associated tracks over the number of nodes that are in the true track. If there are multiple such associated tracks available for a true track, then the completeness score of that true track

is measured by taking the maximum of all the completeness values, that is,

Completeness score of true track
$$i = \max_{j \in J} \frac{b_{ij}}{n_{(k^i)}}$$
. (13)

Here, b_{ij} represents the number of nodes on true track i that are also present in associated track j. The completeness measure also helps one decide track matching or association, if one uses the associated track attaining the maximum completeness value as the representation of the true track. When track association is done, all the following scenarios are possible: a true track matched to multiple associated tracks, an associated track matched to multiple true tracks, and an associated track not mapped to any true tracks.

B. Comparative Performance Evaluation

To apply our algorithm on the test sets, at first, we tune the parameters adopted in our algorithm using the three training datasets for which we know the true VIDs. To summarize, the parameters used in the algorithms are: $\mu=20$ meters, $\beta_{small}=40,\ \beta_{large}=550,\ \alpha=25$ degrees/second, $\tau=300$ seconds, $\gamma=3,000$ meters and $\eta=20$ meters. We use the same settings for all five test sets. After we ran our algorithm on the test sets, the challenge organizers provided us the true track information for the test sets [1], so that we can compute the performance metrics described in the preceding section. We did not change any aspect of our algorithm once we received the true track information of the test sets.

At the end of the data association challenge, the challenge organizers provided us the values of the performance metrics generated by more than ten participating teams, so that we can compare our performance to the best teams. We are not given the identities of these participants nor the specifics of the approaches used by them. Through this comparison, we know that our method performs overall the second best. However, we are not authorized to disclose the performance statistics of other teams. Readers who wish to get this information will need to reach out directly to the NSF ATD Program and make a request.

In this section, in addition to present the performance of our proposed method, we include the performance of a sample algorithm [1], which was made available by the challenge organizers to all participants at the beginning of the competition. We mark our performance values in bold font whenever it comes out as the best (or tied for the best) in a particular category relative to all participating teams, including that of the sample algorithm. This allows us to demonstrate the competitiveness of our method without breaching the confidentiality constraint.

Table II summarizes the count of erroneous tracks representing the number of missing, extra, merged, broken and swapped tracks, produced by our algorithm in five test sets. Inside each parenthesis, it is the corresponding count of erroneous tracks obtained by using the sample baseline algorithm. Here, having a lower count indicates a better performance and zero is the lowest (best) possible value for these counts.

Our algorithm achieves the best performance in four out of five test sets in terms of missed tracks and five out of five cases in terms of extra tracks. On these two error categories, our method outperforms other participating approaches and ranked the first. As the counts of both missed and extra tracks depend on correctly detecting the opening node of a track, we can say that our approach is the most successful in capturing the start of the tracks. In terms of the counts of merged and broken tracks, which, on the other hand, depends on the capability of marking the end node of a track, our approach again produces superior outcomes and becomes the best method in three test sets and remains very close to the best in the other two sets. In the case of swapped tracks, our approach is the best in two test sets and remains competitive in other three sets (we can confirm our method is in fact ranked the second). In all cases, our method outperforms the sample algorithm comprehensively and by a large margin.

In Table III, we summarize the continuity scores and completeness scores generated by our algorithm and the sample algorithm in five test sets. A higher value of the completeness scores indicates a better performance with '1' being the best. A continuity score of '1' indicates that all the true track segments are perfectly captured by the associated track segments.

We further breakdown the completeness score into the mean completeness scores and the median completeness scores. Both completeness scores represent the central tendency of the completeness scores across the numerous tracks in a test set. Because the distribution of the completeness score is not symmetric, the two values do not agree. Their disagreement suggests a skewness towards the high end of the score, i.e., completeness = 1. Moreover, a mean completeness score of '1' indicates that all of the nodes from each true track are perfectly captured by an associated track, and a median completeness score of '1' indicates that more than half of the true tracks are completely captured by the associated tracks. In the category of median completeness scores, our algorithm produces the best result in all five test sets. In the continuity scores and mean completeness scores, our algorithm produces the best result in two test sets and very competitive results in the other three test sets (again we confirm that ours is ranked the second). Also, we again comprehensively outperform the sample algorithm in all categories across all five test sets and our performance is on top with a comfortable margin.

Finally, in Table IV, we show the number of associated tracks identified by our algorithm and the sample algorithm, as compared to the actual number of tracks. Here, our approach identified the number of tracks correctly for all five test cases.

To visualize the track association outcome of our approach, we highlight the tracking performance of our algorithm for five test datasets in Fig. 7. The true tracks are represented by thicker black dots whereas the associated tracks are represented by thinner multicolored dots.

Overall, we can say that our algorithm achieved superior performance across all metrics which are specifically designed to evaluate the performance of a track association approach in different aspects of association (i.e., identifying a new track, marking the end of a track, avoid mixing two different but close tracks that often cross each other, correctly identifying each segment of a track, etc.). We believe, it portrays the strength of our approach in dealing with complicated track

association problems.

C. Performance Comparison with MOTs

We also implemented two popular MOT approaches, namely GNN (Global nearest neighbor) and JPDA (Joint Probabilistic Data Association) to evaluate their performances with respect to our association approach. We summarized their performances in Table V- VII to highlight their competitive edge. Among all the comparison cases, there is not a single instance that these MOT methods outperform our proposed method. In many (albeit not all) cases, they are even worse than the sample algorithm (the baseline method) provided by the data challenge organizers.

According to our observation, these MOTs lose their edge when the vessels are very close to one another, when tracks cross each other and when data collection frequency is high. They severely under-perform in the presence of time gaps and when the number of vessels to track is high (density of vessels in the surveillance area). These factors explain their poor association performance in test sets #3, #4 and #5. In fact, all these characteristics made the challenge problems harder, and thus difficult for the existing algorithms to associate vessels to the true tracks correctly. The data challenge organizers, NGA, in collaboration with NSF, purposely included these tough characteristics in the challenge datasets in order to identify the most competent track association algorithms.

D. Performance Comparison with an Open Source AIS Dataset

To further evaluate the efficacy of our proposed approach, we apply it to an additional open source AIS dataset [50]. The dataset contains an extra variable named 'heading', as compared to the AIS datasets used earlier. To be consistent, we discard this extra variable, so that our existing algorithm can be used without any further modification. To make the dataset size similar to what we have processed so far, we downsample this new dataset by considering first 25,000 instances. The downsampled dataset contains 186 vessels, much more than the number of vessels included in any datasets in the Data Association Challenge. We apply our algorithm to this downsampled dataset and compute the performance metrics. We summarize this performance evaluation in Table VIII, where we also listed the performance of the sample algorithm designed by the competition organizers. Our algorithm achieved the best performance across all metrics and outperformed the sample algorithm by a large margin. We want to mention particularly that our algorithm detected 182 vessels, 2% different from the true number (186), whereas the sample algorithm detected 259 vessels, 40% more than the true number.

VI. CONCLUSION

In this paper, we propose a track association algorithm that associates the AIS data points to true tracks/clusters. Our research found that for the surface vessels, one can reasonably anticipate the next node location based on the AIS measurements taken at the last location of the vessel. Contrasting the

TABLE II: Counts of erroneous tracks produced by our algorithm. Values in the parenthesis are the corresponding count produced by the sample algorithm. Bold indicates the best performance among all participants.

Categories	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
Missed Tracks	0 (1)	0 (1)	0 (7)	2 (7)	1 (7)
Extra Tracks	0 (0)	0 (0)	0 (11)	2 (12)	1 (14)
Merged Tracks	0 (6)	0 (5)	1 (15)	0 (13)	1 (15)
Broken Tracks	0 (5)	0 (4)	1 (91)	0 (18)	1 (22)
Swapped Tracks	0 (729)	0 (1,296)	169 (6,484)	44 (7,953)	1,250 (8,452)

TABLE III: Continuity and completeness score produced by our algorithm. Values in the parenthesis are the corresponding scores produced by the sample algorithm. Bold indicates the best performance among all participants.

Performance Measures	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
Continuity Score	1.000 (0.839)	1.000 (0.935)	0.945 (0.636)	0.988 (0.755)	0.998 (0.699)
Completeness Score (mean values)	1.000 (0.807)	1.000 (0.845)	0.984 (0.604)	0.883 (0.628)	0.916 (0.613)
Completeness Score (median values)	1.000 (0.843)	1.000 (1.000)	1.000 (0.546)	1.000 (0.575)	1.000 (0.648)

TABLE IV: The number of associated tracks recovered by our algorithm and the sample algorithm. Our algorithm recovered perfectly all true track numbers which were unknown to us prior to the application of our method.

Algorithms	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
True Number	8	10	25	25	25
Our Algorithm	8	10	25	25	25
Sample Algorithm	7	9	29	30	32

TABLE V: Counts of missed and broken tracks (separated by a ";") produced by the competing algorithms. Smaller the better. Bold indicates the best performance.

Algorithms	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
Our Algorithm	0;0	0;0	0;0	2;2	1;1
Tracker GNN	3;12	10;179	7;470	11;284	11;174
Tracker JPDA	3;708	10;2233	3; 5775	8;5308	5; 6007
Sample Algorithm	1;0	1;0	7;11	7;12	7;14

anticipated location and the current observed location is an effective strategy to carry out online track association. But a pure trajectory anticipation is not sufficient. In order to have a robust track association, one has to deal with a number of operational complexities. For instance, as the geodesic distance between the anticipated and current locations gets greater, the effectiveness of the anticipation-based association deteriorates. To tackle these complexities, we spatially divide the area under consideration and devise adaptive policies that generate tracks depending on vessel locations. We also turned certain unique characteristics of maritime vessels and their movement into a set of guidelines that ensure a high success rate of the track association process. Using five test datasets provided by the NSF ATD data association challenge, we demonstrate that our algorithm produces superior tracking performance. Our approach also comes superior when compared with two state of art MOT approaches and evaluated on an open source AIS dataset.

The proposed approach is one of the very first attempts to associate trajectory observations to their tracks in real time using a spatio-temporal framework. We are confident that it will advance the trajectory analytics field beyond the regime of the traditional trajectory clustering approaches. The proposed approach can be used for the purpose of dynamic threat detection involving simultaneous movement of multiple objects. In the future study, more than using only the last node location in a number of decision steps, one can consider a small sequence of nodes, or the latent features extracted

from them, to perform the track association process. Doing so could possibly help fix the residual problems associated with the accuracy of node location prediction.

ACKNOWLEDGMENT

Imtiaz Ahmed and Mikyoung Jun acknowledge support by NSF DMS-1925119 under the NSF's Algorithms for Threat Detection (ATD) program. Analysis results are derived from AIS data provided by the US Coast Guard at https://www.navcen.uscg.gov. The data was modified by the National Geospatial-Intelligence Agency in support of the National Science Foundation's ATD program.

REFERENCES

- [1] D. Mercer. Algorithms for threat detection/2019/atd2019. (Date accessed: 10/28/2020). [Online]. Available: https://gitlab.com/algorithms-for-threat-detection/2019/atd2019
- [2] J. H. Ford, D. Peel, D. Kroodsma, B. D. Hardesty, U. Rosebrock, and C. Wilcox, "Detecting suspicious activities at sea based on anomalies in automatic identification systems transmissions," *PloS One*, vol. 13, no. 8, p. e0201640, 2018.
- [3] Z. Ou and J. Zhu, "AIS database powered by GIS technology for maritime safety and security," *Journal of Navigation*, vol. 61, no. 4, pp. 655–665, 2008.
- [4] T. Vincenty, "Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations," *Survey Review*, vol. 23, no. 176, pp. 88–93, 1975.

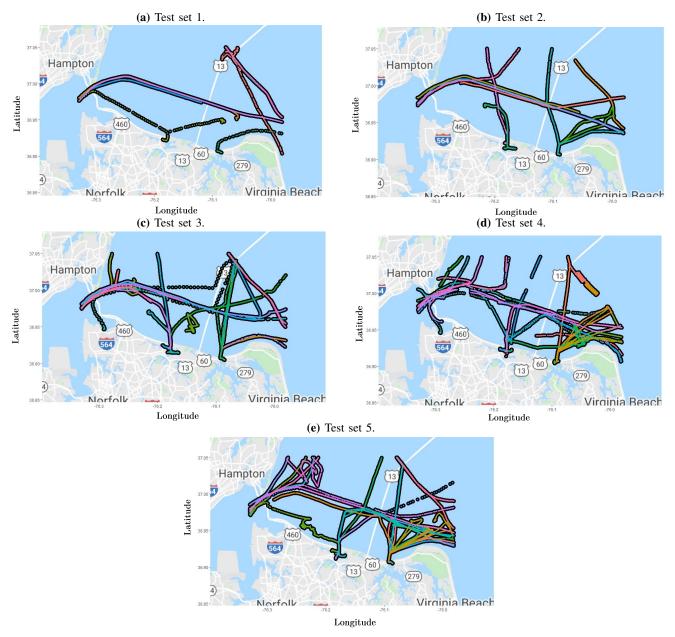


Fig. 7: Tracking performance of our algorithm. Thicker black dots represent the true tracks and thinner multicolored dots represent the associated tracks.

- [5] G. Pallotta, M. Vespe, and K. Bryan, "Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction," *Entropy*, vol. 15, no. 6, pp. 2218–2245, 2013.
- [6] D. Caveney, "Numerical integration for future vehicle path prediction," in 2007 American Control Conference. IEEE, 2007, pp. 3906–3912.
- [7] E. Semerdjiev and L. Mihaylova, "Variable-and fixed-structure augmented interacting multiple model algorithms for manoeuvring ship tracking based on new ship models," *International Journal of Applied Mathematics and Computer Science*, vol. 10, no. 3, pp. 591–604, 2000.
- [8] A. Khan, C. Bil, and K. E. Marion, "Ship motion prediction for launch and recovery of air vehicles," in Proceedings of OCEANS 2005 MTS/IEEE. IEEE, 2005,

- pp. 2795–2801.
- [9] J. Joseph, F. Doshi-Velez, A. S. Huang, and N. Roy, "A Bayesian nonparametric approach to modeling motion patterns," *Autonomous Robots*, vol. 31, no. 4, p. 383, 2011
- [10] G. Pallotta, S. Horn, P. Braca, and K. Bryan, "Contextenhanced vessel prediction based on Ornstein-Uhlenbeck processes using historical AIS traffic patterns: Real-world experimental results," in *17th International Conference on Information Fusion (FUSION)*. IEEE, 2014, pp. 1–7
- [11] C.-H. Yang, C.-H. Wu, J.-C. Shao, Y.-C. Wang, and C.-M. Hsieh, "Ais-based intelligent vessel trajectory prediction using bi-lstm," *IEEE Access*, vol. 10, pp. 24302–24315, 2022.

TABLE VI: Counts of extra, merged, and swapped tracks (separated by ";") produced by the competing algorithms. Smaller the better. Bold indicates the best performance.

Algorithms	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
Our Algorithm	0;0;0	0;0;0	1;1;169	0;0;44	1;1;1250
Tracker GNN	1;10;38	17;179;1670	4;467;3038	7;280;5532	4;167;9887
Tracker JPDA	3;708;554	2;2235;2437	5;5777;4025	8;5308;11030	10;6012;13341
Sample Algorithm	6;5;729	5;4;1296	15;91;6484	13;18;7953	15;22;8452

TABLE VII: Continuity and mean completeness scores (separated by a ";") produced by the competing algorithms. Values between 0 and 1. Greater the better. Bold indicates the best performance.

Algorithms	Test Set-1	Test Set-2	Test Set-3	Test Set-4	Test Set-5
Our Algorithm	1.00;1.00	1.00;1.00	0.95;0.98	0.99;0.88	1.00;0.92
Tracker GNN	0.98;0.98	0.88;0.77	0.65;0.61	0.62;0.49	0.78;0.49
Tracker JPDA	0.46;0.04	0.39;0.01	0.28;0.01	0.10;0.01	0.19;0.01
Sample Algorithm	0.84;0.81	0.94;0.85	0.64;0.60	0.76;0.63	0.70;0.61

TABLE VIII: Performance comparison of our algorithm and the sample algorithm using the additional AIS dataset [50]. Bold indicates the best performance.

Algorithms	Missed Tracks	Extra Tracks	Merged Tracks	Broken Tracks	Swapped Tracks	Continuity Score	Mean Completeness Score	Median Completeness Score
Our Algorithm	38	34	27	23	2,868	0.93	0.93	1
Sample Algorithm	104	177	59	132	15,552	0.35	0.33	0.25

- [12] H. Ma, Y. Zuo, and T. Li, "Vessel navigation behavior analysis and multiple-trajectory prediction model based on ais data," *Journal of Advanced Transportation*, vol. 2022, 2022.
- [13] L. P. Perera and C. G. Soares, "Ocean vessel trajectory estimation and prediction based on extended Kalman filter," in *The Second International Conference on Adaptive and Self-Adaptive Systems and Applications*, 2010, pp. 14–20.
- [14] A. Stateczny and W. Kazimierski, "Multisensor tracking of marine targets: Decentralized fusion of Kalman and neural filters," *International Journal of Electronics and Telecommunications*, vol. 57, pp. 65–70, 2011.
- [15] B. R. Dalsnes, S. Hexeberg, A. L. Flåten, B.-O. H. Eriksen, and E. F. Brekke, "The neighbor course distribution method with Gaussian mixture models for AIS-based vessel trajectory prediction," in 2018 21st International Conference on Information Fusion (FUSION). IEEE, 2018, pp. 580–587.
- [16] E. Tu, G. Zhang, S. Mao, L. Rachmawati, and G.-B. Huang, "Modeling historical AIS data for vessel path prediction: A comprehensive treatment," arXiv preprint arXiv:2001.01592, 2020.
- [17] 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 Transactions on Intelligent Trans*portation Systems, vol. 19, no. 5, pp. 1559–1582, 2017.
- [18] M. P. Muresan, S. Nedevschi, and R. Danescu, "Robust data association using fusion of data-driven and engineered features for real-time pedestrian tracking in thermal images," *Sensors*, vol. 21, no. 23, p. 8005, 2021.
- [19] N. Wang, W. Zhou, J. Wang, and H. Li, "Transformer meets tracker: Exploiting temporal context for robust visual tracking," in *Proceedings of the IEEE/CVF Con-*

- ference on Computer Vision and Pattern Recognition, 2021, pp. 1571–1580.
- [20] G. Yuan, P. Sun, J. Zhao, D. Li, and C. Wang, "A review of moving object trajectory clustering algorithms," *Artificial Intelligence Review*, vol. 47, no. 1, pp. 123–144, 2017.
- [21] K. Toohey and M. Duckham, "Trajectory similarity measures," *Sigspatial Special*, vol. 7, no. 1, pp. 43–50, 2015.
- [22] L. Chen, M. T. Özsu, and V. Oria, "Robust and fast similarity search for moving object trajectories," in *Pro*ceedings of the 2005 ACM SIGMOD International Conference on Management of Data, 2005, pp. 491–502.
- [23] M. Vlachos, M. Hadjieleftheriou, D. Gunopulos, and E. Keogh, "Indexing multidimensional time-series," *The VLDB Journal*, vol. 15, no. 1, pp. 1–20, 2006.
- [24] Z. Zhang, K. Huang, and T. Tan, "Comparison of similarity measures for trajectory clustering in outdoor surveillance scenes," in *18th International Conference on Pattern Recognition (ICPR'06)*, vol. 3. IEEE, 2006, pp. 1135–1138.
- [25] R. C. De Amorim and B. Mirkin, "Minkowski metric, feature weighting and anomalous cluster initializing in K-means clustering," *Pattern Recognition*, vol. 45, no. 3, pp. 1061–1075, 2012.
- [26] T. Zhang, R. Ramakrishnan, and M. Livny, "BIRCH: An efficient data clustering method for very large databases," ACM Sigmod Record, vol. 25, no. 2, pp. 103–114, 1996.
- [27] L. Wang, P. Chen, L. Chen, and J. Mou, "Ship ais trajectory clustering: An hdbscan-based approach," *Journal of Marine Science and Engineering*, vol. 9, no. 6, p. 566, 2021.
- [28] W. Wang, J. Yang, and R. Muntz, "STING: A statistical information grid approach to spatial data mining," in *VLDB*, 1997, pp. 186–195.
- [29] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T.

- Shen, "Discovery of convoys in trajectory databases," arXiv preprint arXiv:1002.0963, 2010.
- [30] C.-C. Hung, W.-C. Peng, and W.-C. Lee, "Clustering and aggregating clues of trajectories for mining trajectory patterns and routes," *The VLDB Journal*, vol. 24, no. 2, pp. 169–192, 2015.
- [31] H. Li, J. Liu, K. Wu, Z. Yang, R. W. Liu, and N. Xiong, "Spatio-temporal vessel trajectory clustering based on data mapping and density," *IEEE Access*, vol. 6, pp. 58 939–58 954, 2018.
- [32] P. Sheng and J. Yin, "Extracting shipping route patterns by trajectory clustering model based on automatic identification system data," *Sustainability*, vol. 10, no. 7, pp. 2327–2340, 2018.
- [33] M. Nanni and D. Pedreschi, "Time-focused clustering of trajectories of moving objects," *Journal of Intelligent Information Systems*, vol. 27, no. 3, pp. 267–289, 2006.
- [34] S. Mitsch, A. Müller, W. Retschitzegger, A. Salfinger, and W. Schwinger, "A survey on clustering techniques for situation awareness," in *Asia-Pacific Web Conference*. Springer, 2013, pp. 815–826.
- [35] G. Yuan, S. Xia, L. Zhang, Y. Zhou, and C. Ji, "An efficient trajectory-clustering algorithm based on an index tree," *Transactions of the Institute of Measurement and Control*, vol. 34, no. 7, pp. 850–861, 2012.
- [36] C. Panagiotakis, N. Pelekis, I. Kopanakis, E. Ramasso, and Y. Theodoridis, "Segmentation and sampling of moving object trajectories based on representativeness," *IEEE Transactions on Knowledge and Data Engineering*, vol. 24, no. 7, pp. 1328–1343, 2011.
- [37] S. Wang, L. Wu, F. Zhou, C. Zheng, and H. Wang, "Group pattern mining on moving objects uncertain trajectories," *International Journal of Computers Communications & Control*, vol. 10, no. 3, pp. 428–440, 2015.
- [38] R. Fileto, A. Raffaetà, A. Roncato, J. A. Sacenti, C. May, and D. Klein, "A semantic model for movement data warehouses," in *Proceedings of the 17th International Workshop on Data Warehousing and OLAP*, 2014, pp. 47–56.
- [39] K. Wolsing, L. Roepert, J. Bauer, and K. Wehrle, "Anomaly detection in maritime ais tracks: A review of recent approaches," *Journal of Marine Science and Engineering*, vol. 10, no. 1, p. 112, 2022.
- [40] P. Fu, H. Wang, K. Liu, X. Hu, and H. Zhang, "Finding abnormal vessel trajectories using feature learning," *IEEE Access*, vol. 5, pp. 7898–7909, 2017.
- [41] C. Chen, D. Zhang, P. S. Castro, N. Li, L. Sun, S. Li, and Z. Wang, "iBOAT: Isolation-based online anomalous trajectory detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 2, pp. 806–818, 2013.
- [42] P.-R. Lei, "A framework for anomaly detection in maritime trajectory behavior," *Knowledge and Information Systems*, vol. 47, no. 1, pp. 189–214, 2016.
- [43] S. Blackman and R. Popoli, Design and Analysis of Modern Tracking Systems. Norwood, MA: Artech House, 1999.

- [44] Y. Bar-Shalom and X.-R. Li, *Multitarget-multisensor Tracking: Principles and Techniques*. YBs Storrs, CT, 1995, vol. 19.
- [45] D. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [46] G. Bishop, G. Welch *et al.*, "An introduction to the kalman filter," *Proceedings of SIGGRAPH*, *Course*, vol. 8, no. 27599-23175, p. 41, 2001.
- [47] C. F. Karney, "Algorithms for geodesics," *Journal of Geodesy*, vol. 87, no. 1, pp. 43–55, 2013.
- [48] H. Goodwin, "The Haversine in nautical astronomy," in *US Naval Institute Proceedings*, vol. 36, no. 3, 1910, pp. 735–746.
- [49] R. J. Hijmans, E. Williams, C. Vennes, and M. R. J. Hijmans, "Package geosphere," *Spherical Trigonometry*, vol. 1, p. 7, 2017.
- [50] "Marine traffic automatic identification system (AIS): Open research maritime vessel tracking dataset," (Date Accessed: 01/16/2022). [Online]. Available: https://www.marinetraffic.com/research/dataset/marinetraffic-automatic-identification-system-ais/



Imtiaz Ahmed received B.Sc. and M.Sc. in Industrial & Production Engineering from Bangladesh University of Engineering & Technology, Dhaka, Bangladesh in 2012 and 2014 respectively. He received his Ph.D. in Industrial Engineering from Texas A&M University in 2020. He is currently working as an Assistant Professor in the Department of Industrial & Management Systems Engineering at West Virginia University. His research interests are in data analytics, machine learning and quality control.



Mikyoung Jun received her Ph.D. in Statistics from University of Chicago (2005). She is currently a Professor and also a ConocoPhillips Data Science Professor of the Department of Mathematics at the University of Houston. One of her main research areas is spatio-temporal modeling for environmental and climate applications, especially development of parametric non-stationary and non-separable covariance functions for univariate, as well as multivariate processes. Dr. Jun is a fellow of ASA, and an elected member of ISI.



Yu Ding (M'01, SM'11) received B.S. from University of Science & Technology of China (1993); M.S. from Tsinghua University, China (1996); M.S. from Penn State University (1998); received Ph.D. in Mechanical Engineering from University of Michigan (2001). He is currently the Mike and Sugar Barnes Professor of Industrial & Systems Engineering and a Professor of Electrical & Computer Engineering at Texas A&M University. His research interests are in quality and data science. Dr. Ding is a fellow of IIE, a fellow of ASME, a senior member of IEEE,

and a member of INFORMS.