Situational Assessment using Indicator Kriging for Fleet Tracking and Prediction

Esther Jose
Industrial and Systems Engineering
University at Buffalo
Buffalo, NY
estherjo@buffalo.edu

Rajan Batta
Industrial and Systems Engineering
University at Buffalo
Buffalo, NY
batta@buffalo.edu

Moises Sudit
Industrial and Systems Engineering
University at Buffalo
Buffalo, NY
sudit@buffalo.edu

***Abstract*—Maritime fleet tracking is a critical piece of naval operations. Leveraging the inherent spatial and temporal autocorrelation of vessels in a fleet, we use spatio-temporal Kriging, an interpolation technique, to estimate the likelihood of finding a vessel at a specific location. This estimation is based solely on the current and/or past locations of other vessels within the fleet. We do this by first fitting covariance models to observed fleet movements. We then use spatio-temporal indicator Kriging to forecast the locations of vessels in a fleet at different times, with or without new information. Our results indicate a notable improvement in accuracy, ranging from 60 to 90% compared to a baseline model. We measure accuracy using ROC AUC values. Furthermore, our study reveals that tracking only a subset of vessels within a fleet significantly enhances understanding of the entire fleet's movements. However, the number of vessels that needs to be tracked increases as we move further from the last observation of the entire fleet. Future extensions of our work include integrating additional situational information, using other spatio-temporal interpolation techniques, and expanding its application beyond maritime fleets.**

***Index Terms*—context-based information fusion, fleet tracking, maritime surveillance, Kriging, spatio-temporal Kriging**

I. Introduction

Central to effective maritime surveillance is the detection and tracking of ships, a task often facilitated by sensor networks that utilize radar, thermal imaging, and low-resolution optical data. By combining these data sources, surveillance capabilities are enhanced, leading to more effective detection and tracking of ships [10]. Many different sources of data and methods of data fusion are employed to both detect and track vessels of interest.

Detection efforts aim to accurately identify vessel or ship presence in water. [21] enhances ship detection through a machine learning approach that incorporates a channel weighting mechanism and spatial information fusion to leverage contextual features effectively. [8] uses an anomaly detection algorithm derived from continuum fusion methods with multispectral satellite data to detect ships. [14] improves ship detection by fusing low-resolution multispectral data with spectral and thermal features to better mask clouds and achieve higher accuracy. [9] proposes a method that combines Automatic Identification System (AIS) data with remote sensing images using feature vectors for ship recognition. [20] applies constant false alarm rate (CFAR) detection on polarimetric Synthetic Aperture Radar (Pol-SAR) data, then fuses the results to improve detection accuracy.

Tracking ships, on the other hand, focuses on predicting the trajectory or movement of vessels or ships. For example, [6] employs a factor graph to track vessels by integrating data from cameras positioned along a harbor. Kalman filters are also commonly used to track or predict the trajectory of ships, using data from GPS [3] or radar [16], among other sources.

Many of these works focus on combining data from different sensors or different types of sensors. In contrast, context-based information fusion involves the use of information not directly about the object or target of interest itself, but about the environment or situation in which the object is situated. For a comprehensive review of context-based information fusion and its applications, readers can refer to [17]. Understanding the context surrounding an entity significantly enhances our comprehension of its existence and behavior. For instance, when tracking a vehicle on a road network, knowledge about the road layout is essential for accurate tracking. Similarly, in a maritime context, contextual information encompasses coastline geography, currents, weather conditions, bathymetry, etc.

Such contextual information has indeed been used to enhance maritime vessel tracking. For example, [13] uses evidence theory to integrate various geographical contexts to improve coastal monitoring. In another example, [5] uses Joint Probabilistic Data Association and Kalman Filters to fuse AIS and High-Frequency Surface-Wave (HFSW) radar data, improving maritime situational awareness over very large areas. Additionally, [19] uses contextual information including bathymetry, nearby port locations, nearby vessels, and knowledge of preferred routes to improve maritime surveillance of vessels or ships of interest. They incorporate these factors into their satellite-extended-vessel traffic service (SEV) data association system, assigning different weights to influences, which can be either 'attractive' or 'repulsive' (e.g., a port is attractive for vessels heading toward it and repulsive for those moving away).

The authors gratefully acknowledge support for this research from the Air Force Office of Scientific Research (AFOSR), as part of the Space University Research Initiative (SURI), grant FA9550-22- 1-0092 (grant co-principal investigator: M. Sudit from University at Buffalo, The State University of New York).

Fleet tracking involves monitoring multiple vessels or ships that are part of a fleet. While the sensor data and contextual information discussed earlier are indeed valuable for enhancing the tracking of individual vessels, they do not specifically address the unique dynamics of fleets. To our knowledge, a crucial contextual factor that has not been adequately integrated into fleet tracking is the spatial autocorrelation among vessels within fleets. Vessels within fleets naturally exhibit both spatial and temporal autocorrelation. This means that if one vessel from a fleet is detected in a certain location, there's a higher probability of another vessel from the same fleet being nearby. Additionally, vessels within a fleet are likely to be found near where they were previously observed. The extent of this autocorrelation and the distance over which vessels within a fleet influence each other can vary depending on the specific circumstances.

Kriging, a geostatistical interpolation technique, is widely applied in spatial analysis to estimate values at unmeasured locations by considering data from neighboring locations, thus incorporating spatial autocorrelation information [7]. Using such techniques presents a promising approach for enhancing fleet tracking performance. It is worth noting here that one of the factors used by [19] to track a vessel is its proximity to other vessels, albeit not necessarily from the same fleet. The underlying principle here is that vessels do not get very close to other vessels, in order to to avoid collision. Thus, it's improbable for a vessel to be in close proximity to other vessels. On the other hand, we know that vessels within a fleet tend to travel near each other, and we leveraging this knowledge to enhance our ability to track fleets effectively. In contrast, [19] focuses on tracking individual vessels.

In this paper, we address this gap by introducing a novel framework that incorporates spatial and temporal autocorrelation into fleet tracking through spatio-temporal Kriging. Specifically, we use spatio-temporal indicator Kriging to predict the probable locations of vessels within a given fleet, based on the observed locations of vessels from that fleet at different times. Our aproach aligns with Level 2 of the JDL (Joint Directors of Laboratories) information fusion model, called Situation Assessment. Situation assessment can be defined as the "estimation and prediction of entity states on the basis of inferred relations among entities" [17], [18]. In our context, these relationships among entities (or vessels) are represented by spatial and temporal autocorrelation, which we utilize to determine probable vessel locations in a fleet.

The contributions of our work are as follows:

1) Development of a framework utilizing Kriging for contextual information fusion in maritime fleet surveillance.
2) Demonstration of the effectiveness of our method through an illustrative example.

The remainder of this paper is structured as follows: In Section II, we explain spatio-temporal indicator Kriging and Receiver Operating Characteristic (ROC), which we use to assess the accuracy of our predictions. In Section III, we explain how we generate data for our example, fit our model to the generated data, and present the results. Finally, we discuss the limitations of this work and present future extensions in Section IV.

II. Methodology

In Section II-A, we explain how Kriging uses covariance structures to interpolate the value of a variable of interest at an unsampled location. We also explain different covariance models that can be used to represent the covariance structure of points across space and time. In Section II-B, the Receiver Operating Characteristic is detailed, which is used to assess the validity of our predictions later on in Section III.

A. Spatio-temporal Indicator Kriging

Kriging: Kriging is a geostatistical interpolation method used to estimate values at unsampled locations within a region of interest. It typically utilizes the spatial autocorrelation between sampled points to make predictions. The main idea behind Kriging is to model the spatial dependence structure of the variable of interest using a variogram. The variogram quantifies the spatial variability of the variable as a function of distance and provides information about the spatial autocorrelation of the data it is fit to.

The variogram $\gamma(h)$ is defined as:

$$\gamma(h) = \frac{1}{2N(h)} \sum_{i=1}^{N(h)} \left(Z(s_i + h) - Z(s_i)\right)^2 \tag{1}$$

where h is the lag distance, $N(h)$ is the number of pairs of points separated by the lag distance h, and $Z(s_i)$ is the value at sampled location s_i. The variogram has three important characteristics: the range, the sill, and the nugget [2]. The sill is the peak variance between two points in a variogram. Since the variance between two points increases the farther away they are from each other, the variance increases with distance. Then, at a certain distance known as the range, the variogram stabilizes or plateaus. The total variance where the variogram levels off is referred to as the sill. Theoretically, at a lag or a separation distance of $h = 0$, the variance given by $\gamma(h)$ should be zero. However, in practice, a small variance is often observed when $h = 0$, called the nugget. That is, when $h = 0$, $\gamma(h)$ is equal to the nugget value.

Now, the Kriging estimator at an unsampled location s_0 can be expressed as:

$$\hat{Z}(s_0) = \sum_{i=1}^{n} \lambda_i Z(s_i) \tag{2}$$

where $\hat{Z}(s_0)$ is the estimated value at unsampled location s_0, while $Z(s_i)$ is the value at the sampled location s_i. The Kriging weights, λ_i, are determined from the covariance structure or variogram. For a more comprehensive understanding of Kriging, the reader is referred to [15].

Spatio-temporal Kriging: Spatio-temporal Kriging extends the concept of Kriging to incorporate both spatial and temporal dependencies. It considers the spatial correlation between locations as well as the temporal correlation between time points. This is particularly useful when dealing with datasets that

exhibit both spatial and temporal trends, such as environmental monitoring data or climate data.

The spatio-temporal Kriging estimator at an unsampled location s_0 at time t_0 can be written as:

$$\hat{Z}(s_0, t_0) = \sum_{i=1}^{n} \sum_{j=1}^{m} \lambda_{ij} Z(s_i, t_j) \tag{3}$$

where $\hat{Z}(s_0, t_0)$ is the estimated value at location s_0 and time t_0, $Z(s_i, t_j)$ is the value at the sampled location s_i at time t_j, and λ_{ij} are the spatio-temporal Kriging weights.

Indicator Kriging: Indicator Kriging is a variant of Kriging that is used when dealing with categorical or binary data, where the variable of interest takes on discrete values or represents the occurrence of an event. Then, when estimating the value of the variable at an unsampled location and time, we are given the probability of that event occurring at that specific location and time. The spatio-temporal indicator Kriging estimator for a specific event k at an unsampled location s_0 and time t_0 can be written as:

$$\hat{P}(Z(s_0, t_0) = k) = \sum_{i=1}^{n} \sum_{j=1}^{m} \lambda_{ij} I(Z(s_i, t_j) = k) \tag{4}$$

where $\hat{P}(Z(s_0, t_0) = k)$ is the probability of event k occurring at location s_0 and time t_0, $I(Z(s_i, t_j) = k)$ is an indicator function that equals 1 if the event occurs at location s_i and time t_j, and λ_{ij} are the spatio-temporal indicator Kriging weights.

Covariance models: Kriging estimations heavily rely on variograms, as discussed above, to define the spatial and temporal dependencies within the data. Covariance functions are closely related to variograms. While variograms quantify spatio-temporal variability as a function of distance and time, covariance functions quantify covariance as a function of distance and time. That is, given spatio-temporally autocorrelated data, spatio-temporal variability increases with distance and time, while covariance decreases with distance and time. As defined by [1], the spatial covariance function $C(s_1, s_2)$ between two points in space, s_1 and s_2, is:

$$C(s_1, s_2) = \text{sill} - \gamma(s_1, s_2) \tag{5}$$

This relationship in Equation 5 also holds for spatio-temporal distances. Finding the right covariance function is an important step in using spatio-temporal Kriging. Different spatio-temporal covariance models, such as separable models or metric models, offer different approaches to capture spatial, temporal, and spatio-temporal dependencies. Given a distance h and time τ between any two points in space and time, the different spatio-temporal covariance models used in this paper are defined as in [11]:

1) Metric Model: In this model, the two dimensional spatial distance is simply extended into a three dimensional spatio-temporal distance. To make the distance isotropic, a spatio-temporal anistropic correction parameter, κ, is used. Then, spatial, temporal, and spatio-temporal distances are all considered equally, leading to a joint covariance function, defined as C_j:

$$C_m(h, \tau) = C_j\left(\sqrt{h^2 + (\kappa\tau)^2}\right) \tag{6}$$

2) Sum Metric Model: The sum metric model is an extension of the metric model. In addition to the joint covariance function seen in 6, we also see separate spatial and temporal components:

$$C_{sm}(h, \tau) = C_s(h) + C_t(\tau) + C_j\left(\sqrt{h^2 + (\kappa\tau)^2}\right) \tag{7}$$

where $C_s(h)$ and $C_t(\tau)$ are spatial and temporal covariance functions respectively. In this model, the spatial, temporal, and joint components have separate nugget effects.

3) Simple Sum Metric Model: This model is a more restricted version of the sum metric model. While the covariance function looks the same, it assumes that the spatial, temporal, and joint variograms do not have nuggets. Instead, a single spatio-temporal nugget is used.

4) Separable Model: This model assumes that the covariance between two spatio-temporal points can be factorized into separate spatial and temporal components, represented as:

$$C_{sep}(h, \tau) = C_s(h) C_t(\tau)$$

5) Product Sum Model: This model is an extension to the separable model, and offers additional flexibility:

$$C_{ps}(h, \tau) = k C_s(h) C_t(\tau) + C_s(h) + C_t(\tau)$$

where k is a non-negative real coefficient which is estimated based on the spatial, temporal, and spatio-temporal sill values.

B. Receiver Operating Characteristic

In this paper, we use spatio-temporal indicator Kring to predict the probability that at a given location and time, there exists a ship from a certain fleet. To assess the quality of our predictions, we use a Reciever Operating Characteristic (ROC) curve. The ROC curve and the Area Under the ROC Curve (ROC AUC) are commonly used to evaluate the performance of classification models [4], [12].

The ROC curve is a graphical representation of the performance of a binary classifier across different discrimination thresholds. The advantage of this method is that a specific threshold does not need to be chosen in order to assess the performance of the binary classifer. The curve plots the True Positive Rate (TPR) against the False Positive Rate (FPR) at various threshold settings. TPR represents the proportion of positives that are correctly identified, while FPR represents the proportion of false positives among all negatives. The equation for TPR (Sensitivity) and FPR (1 - Specificity) are as follows:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{8}$$

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \tag{9}$$

where TP is the number of true positives, FN is the number of false negatives, FP is the number of false positives, and TN is the number of true negatives. The ROC AUC represents the area under the ROC curve and provides a single scalar value to summarize the performance of the classifier. It ranges from 0 to 1, where a value closer to 1 indicates better performance. For a binary classifier, a random predictor, the typical benchmark, has an ROC AUC value of 0.5.

III. Example

In this section, we outline our process for generating synthetic data that replicates the movements of spatially autocorrelated ships within a fleet (Section III-A). We then delve into the outcomes of fitting various covariance models to our dataset, along with the eventual model selection (Section III-B). Finally, in Section III-C, we employ the chosen covariance model in spatio-temporal indicator Kriging to predict probable ship locations. We evaluate the prediction accuracy using ROC AUC values.

A. Generation of Synthetic Ship Data

We begin by generating (x, y) positions for 20 ships over 1000 timesteps. Each ship's initial position is randomly determined within the range of 0 to 1 for both x and y coordinates. This ensures that while all ships start from a similar vicinity, their starting positions vary, thereby diversifying the initial conditions for subsequent movements. Subsequently, at each time step of the simulation, the positions of the ships are updated to reflect their movements. This updating process involves two key factors.

Firstly, in order to simulate spatial autocorrelation, a directional bias is introduced in the form of a unit vector. In our scenario, all ships move with a bias of $(1, 0)$, indicating movement to the right. Secondly, random changes in positions are independently calculated for each ship at every timestep, with x and y values ranging between -0.5 and 0.5. This simulates unknown decisions regarding the ships' movements, as well as the inherent stochastic nature of ship movements. Both the directional bias and random changes are scaled to match a given step distance, δ, ensuring that each ship moves at a constant speed. in our scenario, we set the step distance to be $\delta = 0.1$. Given a step distance δ and a unit directional bias of $\mathbf{d}_u$, the directional bias used to update the ship positions, $\mathbf{d}$, is calculated as:

$$\mathbf{d} = \delta \cdot \frac{\mathbf{d_u}}{\|\mathbf{d_u}\|_2} = \delta \cdot \mathbf{d_u} \tag{10}$$

Similarly, the random vector, $\mathbf{r_v}$, is scaled using δ:

$$\mathbf{r} = \delta \cdot \frac{\mathbf{r_v}}{\|\mathbf{r_v}\|_2} \tag{11}$$

The ship movements are updated using a weighted combination of the directional bias and random changes. The weight assigned to the directional bias determines the extent of autocorrelation in the ships' movements, while the weight assigned to the random changes affects the dispersion of the ships from each other. We assign a weight of 0.9 to the directional bias, and subsequently, a weight of 0.1 to the random changes, as we want to create a spatially autocorrelated fleet. The movement of each ship in each timestep after the first one can be calculated as follows, where $\mathbf{P}_t$ is the (x, y) position of the ship at time t, w_d refers to the weight assigned to the directional bias, $\mathbf{d}$ is the directional bias vector, and $\mathbf{r}$ is the vector with a scaled random movement:

$$\mathbf{P}_t = \mathbf{P}_{t-1} + w \cdot \mathbf{d} + (1 - w) \cdot \mathbf{r} \tag{12}$$

In Fig. 1, positions for the twenty simulated ships are shown at different timeslices during the beginning, middle, and end of the 1000-timestep duration.

B. Fitting training data to covariance models

To fit the spatio-temporal data generated in Section III-A to the different covariance models from Section II-A, we construct gridded data in both space and time. In this setup, the value of a cell at a particular timestep is set to 1 if any ships are present in that cell at that time, and 0 otherwise. To create a sufficient training dataset, we monitor ship positions from timestep 120 to 180, within an x range of 11 to 17, and a y range of 0.05 to 1.25. These x and y ranges are determined based on the ship locations during the selected time duration. The grid is configured such that each cell has a width and height of 0.25 units. Moreover, we capture data from every third timeslice within our chosen duration. This approach yields a dataset comprising 3150 observations, as illustrated in Fig. 2. To utilize the *gstat* package in R for fitting spatio-temporal data, we need to establish units for both space and time. We define each timestep as equivalent to 30 seconds and each unit of space as one meter. Consequently, our dataset covers a duration of 30 minutes. It's important to note that setting each timestep to 30 minutes would yield identical results, albeit over a different time scale.

We create a sample or empirical variogram from the training data, and then fit the empirical variogram to each of the five covariance models detailed in Section II-A. To identify the optimal parameters for each covariance model, we employ the L-BFGS-B method [22]. The L-BFGS-B method is an optimization algorithm commonly used for problems where the objective function is smooth and possibly nonlinear, and where some or all of the variables are subject to bounds. It belongs to the family of quasi-Newton methods, which aim to iteratively improve an initial guess for the optimal solution by approximating the Hessian matrix (which represents the second-order derivative of the objective function). In the context of fitting covariance models to empirical variograms, L-BFGS-B iteratively adjusts the parameters of the covariance models within predefined bounds to minimize the Mean Squared Error (MSE) between the fitted theoretical variogram and the empirical variogram [11]. This process

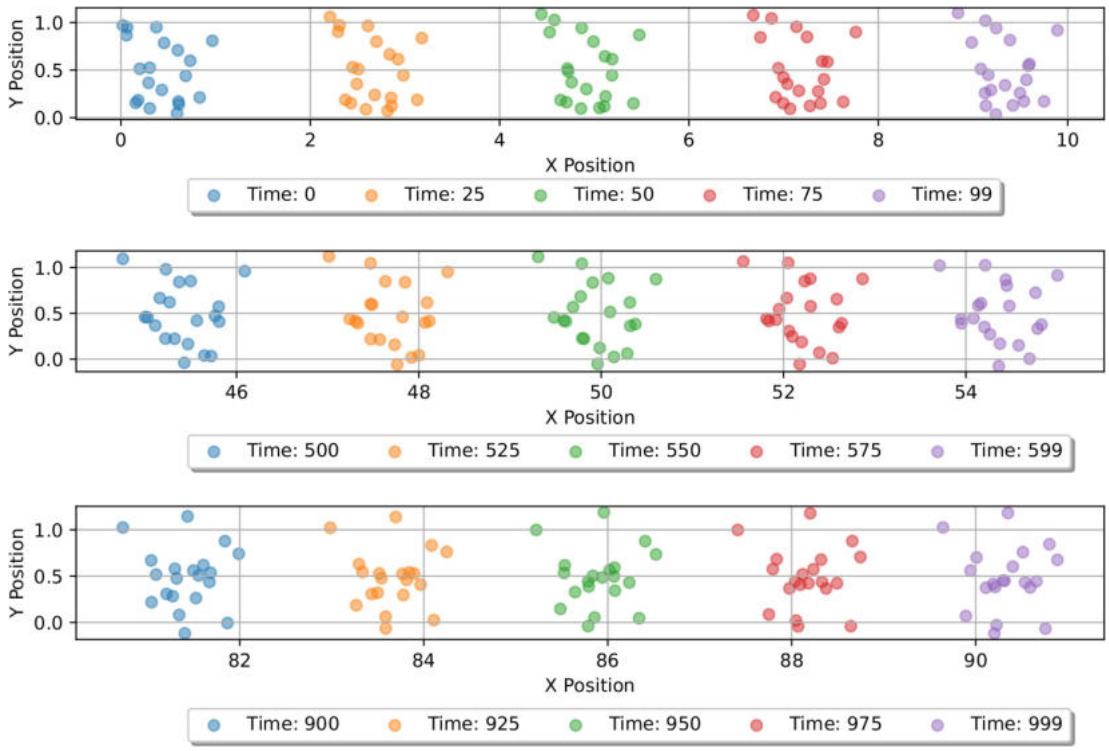

Fig. 1: Ship positions at different time slices over the 1000 timesteps

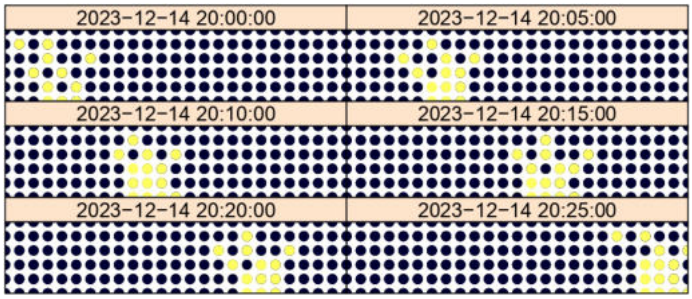

Fig. 2: The gridded training data, with cells colored yellow indicating the presence of a ship, and black indicating no ship presence.

helps in finding the optimal parameters that best capture the spatio-temporal dependence structure observed in the data. The fitted theoretical variograms corresponding to each model are illustrated in Fig. 3. The final MSE values for each model are also presented in Table I.

TABLE I: The MSE values associated with the fit of each covariance model to the sample variogram

Model	MSE
Metric	1.04E-04
Sum Metric	8.85E-05
Simple Sum Metric	1.44E-04
Separable	1.62E-04
Product Sum	1.01E-04

We observe in Table I that the MSE values indicate a relatively small range, varying from 8.85E-05 (for the sum metric) to 1.62E-04 (for the separable model). However, upon examining Fig. 3, we see that the metric, sum metric, and product sum models, despite having lower MSE values compared to the other two models, fail to capture specific patterns present in the empirical variogram. Notably, the empirical variogram displays minimal variance between points in close proximity in both space and time, a characteristic not reflected in these theoretical variograms. This discrepancy suggests that these models exhibit a tendency to overfit, resulting in a flattened average variogram across both space and time. On the other hand, the simple sum metric and separable models demonstrate a better alignment with these observed patterns in the empirical variogram. Consequently, in Section III-C, we experimentally compare the simple sum metric and separable models for predicting ship locations.

C. Prediction

In this section, we predict the probable locations of ships belonging to the fleet. Predictions are generated for timeslices occurring after the latest observed timeslice or timestep in the training data, with and without the inclusion of 'new information'. Here, 'new information' refers to the provision of the true locations of a few ships during the timeslice for which predictions are being made. If we have no new information, we forecast the probability of encountering a ship based only on observed data from 20:00:00 to 20:30:00, which is also our training data as depicted in Fig. 2.

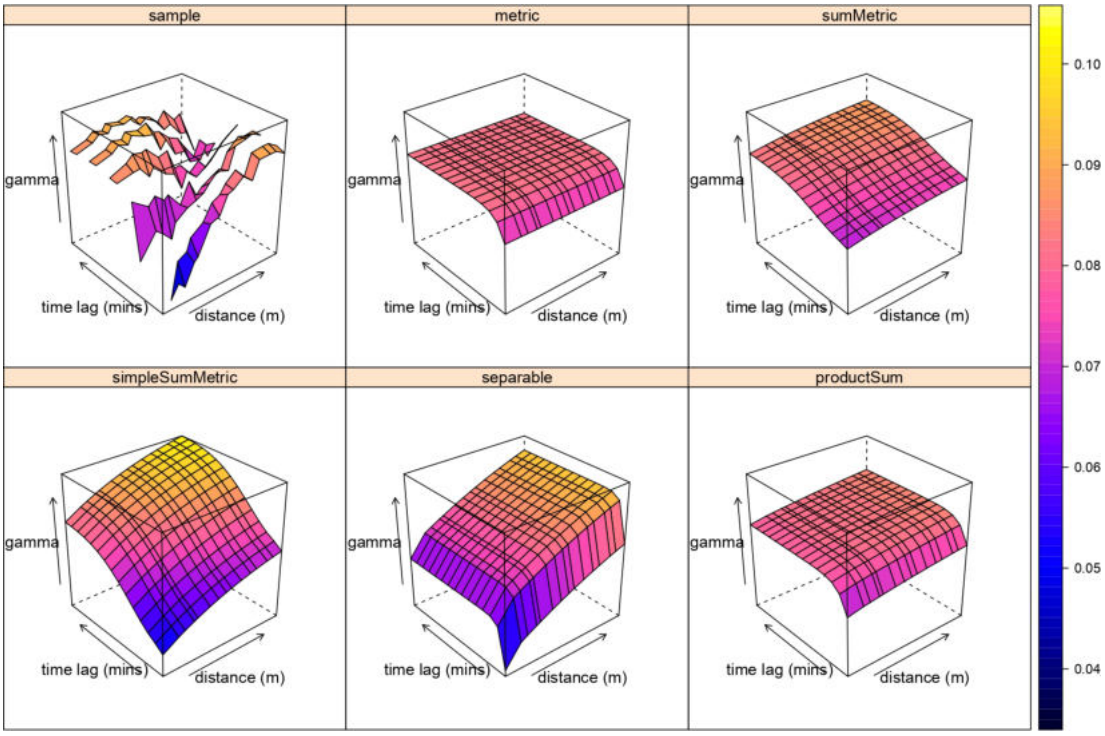

Fig. 3: The empirical and theoretical variograms of the training data

We initially compare the prediction accuracies of separable and simple sum metric covariance models across four different scenarios in Section III-C1, finding the separable model to perform better. Then, in Section III-C2, we conduct experiments to determine how many ships need to be tracked to enhance fleet tracking. Overall, we make predictions over both short ranges such as 1 minute after the last observation (20:31:00) and long ranges such as 2 hours after the last observation (22:30:00), and are able to obtain ROC AUC values of greater than 80% in both of these scenarios.

1) Comparing separable and simple sum metric covariance models: To compare the performance of the separable and simple sum metric models, we conduct forecasts 1 minute or 5 minutes after the last observation, with and without new information. This new information consists of the true locations of five randomly selected ships at the time of prediction. In each of these four scenarios, we utilize either model to predict the probability of a ship existing in a square cell measuring 0.05 in both height and width. The prediction is carried out within a spatial domain spanning an x range of 15 to 19 and a y range of 0.05 to 1.15. We choose this spatial range as it represents a broader area where the fleet could potentially be located, as it accommodates the fleet's bounds of x from 15.8 to 17.2 and y from 0.08 to 1.08 at the last observed time step of 20:30:00. This is assuming that a decision maker is aware the fleet tends to move to the right. Then, the baseline model assumes that a ship from the fleet could be anywhere in this spatial domain, corresponding to an ROC AUC value of 0.5. On the other hand, the ROC AUC values for all four scenarios, using either a separable or a simple sum metric model in the same spatial domain, are presented in Table I. We observe that the separable model outperforms the simple sum metric model in most scenarios.

TABLE II: ROC AUC values for four different prediction scenarios with separable or simple sum metric models

Min. after last obs.	**New info**	**ROC AUC**	
		Separable	***SSM***
1	Yes	0.888	0.751
1	No	0.864	0.551
5	Yes	0.894	0.504
5	No	0.686	0.782

For the cases where we forecast 5 minutes after the last observation, we also present the results graphically, as detailed below. In Fig. 4, we display prediction heatmaps when no new information is given. Here, the simple sum metric model outperforms the separable model. In the scenario with new information, the model has access to the locations of five random ships out of twenty ships in the fleet. The locations of these five ships are shown in Fig. 5, where the dotted blue box indicates the area of prediction. In Fig. 6, we present the prediction heatmaps after the model is given the true locations of these five ships.

We observe that the simple sum metric model displays stronger temporal autocorrelation than the separable model. When given new information 5 minutes later, the separable model, as shown in Fig. 6b, appears to predominantly rely on this new information, disregarding any previous data. Consequently, it produces a heatmap significantly different

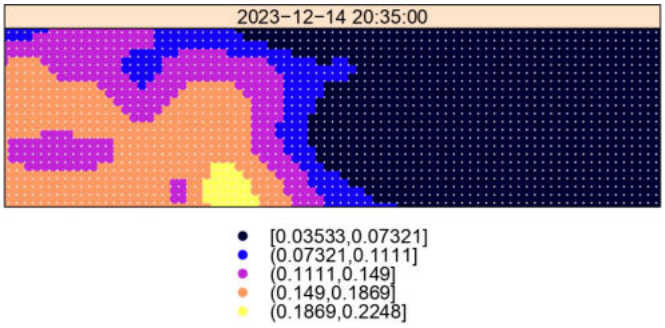

(a) Simple Sum Metric model

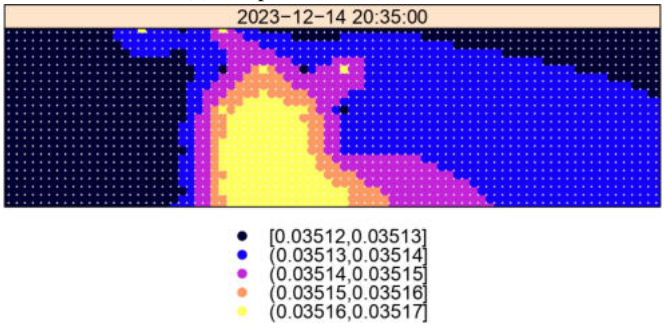

(b) Separable model

Fig. 4: Prediction heatmaps for 5 minutes after the last observation, with no new information

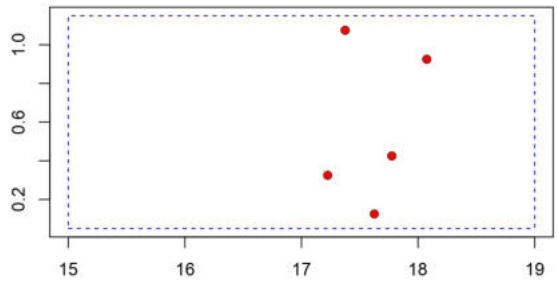

Fig. 5: Locations of 5 ships, 5 minutes after the last observation

from Fig. 4b. On the other hand, the simple sum metric model, in Fig. 6a, considers both the previous data, as well as the new information. While the appearance of its heatmap changes from Fig. 4a, the changes primarily occur in regions influenced by the new information.

Overall, incorporating new information significantly increases prediction accuracy, especially as we move farther away from the last observation. We further investigate the value of new information in the following section. We also notice that, given new information, the separable model outperforms the simple sum metric model. Even without new information, the separable model still performs better when forecasting one minute after the last observation, but not 5 minutes after. Consequently, we opt for the separable covariance model for our future predictions, while acknowledging the potential value of the simple sum metric model for forecasting farther from the last observation when new information is unavailable.

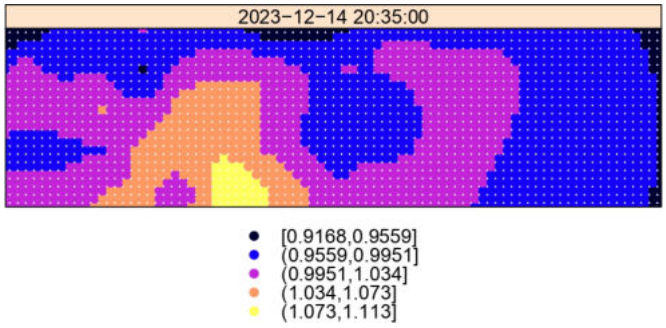

(a) Simple Sum Metric model

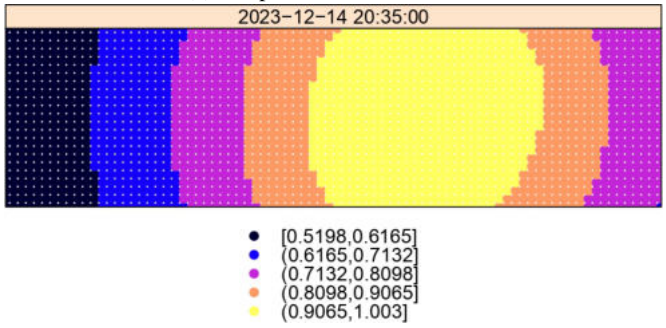

(b) Separable model

Fig. 6: Prediction heatmaps for 5 minutes after the last observation, with new information

2) Amount of new information needed: Tracking the movements of ships is expensive, requiring a balance between monitoring a sufficient number of vessels or ships to estimate the fleet's probable location while mitigating the expenses associated with tracking each additional ship. To understand the impact of the number of ships we track, we make forecasts for 1 minute, 5 minutes, 12.5 minutes, 1 hour, and 2 hours after the last observation. For each forecast, we provide the model with the actual locations of 0 to 5 randomly selected ships. The resulting ROC AUC values for these scenarios are depicted in Fig. 7.

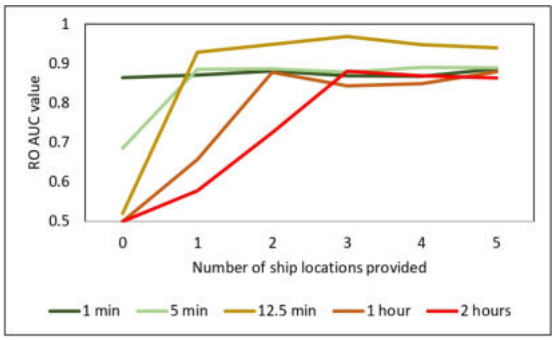

Fig. 7: ROC AUC values for forecasts at different times with different amounts of new information

From this plot, two key observations emerge. First, the need

for new information intensifies as we forecast further from the last observation. As our predictions extend to hours after the last observation, including the locations of two or three ships yields sufficiently accurate predictions. Second, when making predictions hours after the last observation, the absence of new information results in a ROC AUC value of 0.5, akin to that of a random binary classifier in the given space.

IV. Conclusion

In this paper, we highlight the significance of leveraging spatio-temporal autocorrelation structures within a fleet of ships to enhance fleet tracking accuracy. By employing spatio-temporal Kriging, we predict the fleet's locations at various time intervals following the last observed ship locations, with or without new information, achieving ROC AUC values ranging from 0.8 to 0.95. Without utilizing spatio-temporal autocorrelation, decision-makers would only be able to provide approximate bounds on the fleet's likely whereabouts based on prior observations and tracked ship locations, resulting in a ROC AUC value of 0.5. Our model demonstrates a substantial improvement of 60 to 90% over this baseline. We also demonstrate that tracking just a few ships in a fleet can be used to significantly increase the accuracy of predicting the fleet's location.

Our approach forms just one facet of effective fleet tracking and cannot operate in isolation for comprehensive fleet surveillance. Moving forward, we intend to expand our research by incorporating additional contextual information such as bathymetry, prior knowledge, and remote sensing data. Additionally, we will explore the use of other spatio-temporal interpolation methods such as copulas, which can better capture broader movement patterns of a fleet. We also aim to generalize our model to apply to other autocorrelated targets of interest, such as space assets.

References

[1] ArcGIS. Semivariogram and covariance functions. Retrieved February, 2024 from https://desktop.arcgis.com.

[2] ArcGIS. Understanding a semivariogram: The range, sill, and nugget. Retrieved February, 2024 from https://pro.arcgis.com/en/pro-app/3.1/help/analysis/geostatistical-analyst/understanding-a-semivariogram-the-range-sill-and-nugget.htm.

[3] Mansour H Assaf, Voicu Groza, and EMIL M Petriu. The use of kalman filter techniques for ship track estimation. *WSEAS Trans. Syst*, 19:7–13, 2020.

[4] Alex J Bowers and Xiaoliang Zhou. Receiver operating characteristic (roc) area under the curve (auc): A diagnostic measure for evaluating the accuracy of predictors of education outcomes. *Journal of Education for Students Placed at Risk (JESPAR)*, 24(1):20–46, 2019.

[5] Paolo Braca, Raffaele Grasso, Michele Vespe, Salvatore Maresca, and Jochen Horstmann. Application of the jpda-ukf to hfsw radars for maritime situational awareness. In *2012 15th International Conference on Information Fusion*, pages 2585–2592, 2012.

[6] Francesco Castaldo and Francesco AN Palmieri. Data fusion using a factor graph for ship tracking in harbour scenarios. In *Recent Advances of Neural Network Models and Applications: Proceedings of the 23rd Workshop of the Italian Neural Networks Society (SIREN), May 23-25, Vietri sul Mare, Salerno, Italy*, pages 189–196. Springer, 2014.

[7] Noel Cressie. The origins of kriging. *Mathematical geology*, 22:239–252, 1990.

[8] Brian J Daniel, Alan P Schaum, Eric C Allman, Robert A Leathers, and Trijntje V Downes. Automatic ship detection from commercial multispectral satellite imagery. In *Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XIX*, volume 8743, pages 316–323. SPIE, 2013.

[9] Guanglei Fan, Xin Song, Yong Zhao, Zhenguo Yan, and Xueyang Wang. Ship fusion recognition based on ais data and remote sensing image. In *Proceedings of the 2022 6th International Conference on Computer Science and Artificial Intelligence*, pages 116–120, 2022.

[10] Ahmed Gad and M Farooq. Data fusion architecture for maritime surveillance. In *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002.(IEEE Cat. No. 02EX5997)*, volume 1, pages 448–455. IEEE, 2002.

[11] Benedikt Gräler, Edzer J Pebesma, and Gerard BM Heuvelink. Spatio-temporal interpolation using gstat. *R J.*, 8(1):204, 2016.

[12] Zhe Hui Hoo, Jane Candlish, and Dawn Teare. What is an roc curve?, 2017.

[13] Alexandre Jouan, Yannick Allard, and Yves Marcoz. Coastal activity monitoring with evidential fusion of contextual attributes from multi-pass radarsat-1 data. In *7th International Command and Control Research and Technology Symposium, Quebec City*, 2002.

[14] Yong Liu, Libo Yao, Wei Xiong, and Zhimin Zhou. Fusion detection of ship targets in low resolution multi-spectral images. In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 6545–6548. IEEE, 2016.

[15] Margaret A Oliver, Richard Webster, et al. Basic steps in geostatistics: the variogram and kriging. Technical report, Springer, 2015.

[16] Lokukaluge P Perera, Carlos Guedes Soares, et al. Ocean vessel trajectory estimation and prediction based on extended kalman filter. In *The Second International Conference on Adaptive and Self-Adaptive Systems and Applications*, pages 14–20. Citeseer, 2010.

[17] Lauro Snidaro, Jesús García, and James Llinas. Context-based information fusion: A survey and discussion. *Information Fusion*, 25:16–31, 2015.

[18] Alan N Steinberg and Christopher L Bowman. Revisions to the jdl data fusion model. In *Handbook of multisensor data fusion*, pages 65–88. CRC press, 2017.

[19] Michele Vespe, Massimo Sciotti, Fabrizio Burro, Giulia Battistello, and Stefano Sorge. Maritime multi-sensor data association based on geographic and navigational knowledge. In *2008 IEEE Radar Conference*, pages 1–6, 2008.

[20] Wenguang Wang, Yu Ji, and Xiaoxia Lin. A novel fusion-based ship detection method from pol-sar images. *Sensors*, 15(10):25072–25089, 2015.

[21] Quancheng Zhou, Fei Song, Zhenghao Chen, Rui Zhang, Ping Jiang, and Tao Lei. Harbor ship detection based on channel weighting and spatial information fusion. In *Journal of Physics: Conference Series*, volume 1738, page 012057. IOP Publishing, 2021.

[22] Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on mathematical software (TOMS)*, 23(4):550–560, 1997.