

Research Objective

This study develops a satellite-based machine-learning framework for detecting, characterizing, and forecasting maritime vessel congestion along strategically important oil-shipping routes. The analysis focuses on maritime activity surrounding the Strait of Hormuz and nearby shipping regions, where vessel clustering and changes in traffic density may provide information about disruptions or changes in oil and logistics activity. The broader objective is to transform satellite observations into structured time-series signals that can eventually be incorporated into market or logistical analysis for trading signals.

The proposed framework combines Synthetic Aperture Radar (SAR) satellite imagery, spatial feature engineering, convolutional neural networks (CNNs), and temporal modeling. Rather than relying solely on conventional vessel-tracking information, the approach extracts maritime activity directly from satellite imagery and converts the detected vessel distribution into representations suitable for machine-learning analysis.

Satellite Data Acquisition

The first stage of the methodology obtains Sentinel-1 SAR imagery through Google Earth Engine. SAR imagery is particularly useful for maritime analysis because vessels can appear as high-backscatter objects relative to the surrounding ocean surface.

The implementation focuses on a geographic region surrounding the Strait of Hormuz. Sentinel-1 Ground Range Detected imagery is filtered according to geographic coverage, acquisition dates, Interferometric Wide Swath mode, and VV polarization. Only satellite observations covering at least 95% of the selected study region are retained.

The SAR observations are subsequently converted into numerical arrays so that the satellite images can be processed using standard Python-based image-processing and machine-learning techniques.

Vessel Detection and Maritime Congestion Measurement

The second stage converts the SAR observations into estimates of maritime activity. Vessel-like objects are identified by detecting pixels with unusually high radar backscatter relative to the surrounding image.

For each SAR observation, a threshold corresponding to the 99.8th percentile of radar backscatter intensity is calculated. Pixels exceeding this threshold are classified as potential vessel observations. Connected-component analysis is then used to group adjacent bright pixels into individual objects, while a size filter removes objects that are either too small or too large to represent plausible ships.

The resulting vessel detections are aggregated by acquisition date to construct a ship-congestion time series. For the sample shown in the project, 15 usable satellite acquisitions were identified between April 1 and May 13, 2026, and estimated ship-candidate counts ranged from approximately 149 to 172 vessels per image.

Conceptually, the transformation can be written as:

$$SAR\ Image_t \rightarrow Ship\ Detection_t \rightarrow N_t$$

where N_t represents the number of detected ship candidates at observation date t .

Spatial Decomposition of Maritime Traffic

A simple total vessel count does not fully characterize maritime congestion because two observations can contain the same number of ships while exhibiting substantially different spatial distributions. The methodology therefore introduces an additional spatial layer.

The study region is divided into four maritime zones:

- North Lane
- South Lane
- Khasab Bay
- Open Gulf West

Each detected vessel is assigned to one of these geographic zones based on its estimated centroid. The resulting observation for time t can therefore be represented as a regional state vector:

$$X_t = [N_{North,t}, N_{South,t}, N_{Khasab,t}, N_{West,t}]^T$$

This representation captures not only the total number of ships but also the spatial distribution of maritime congestion.

Flow-Network Representation

The regional vessel counts are subsequently transformed into graphical flow-network images. Each geographic region becomes a node, while the relative size of each node represents the number of ships detected within that region.

The purpose of this transformation is to encode the maritime congestion structure in a format that can be interpreted by a convolutional neural network. Rather than asking the model to infer geography directly from raw SAR pixels, the network image provides a structured representation of how maritime activity is distributed across the study area.

$$X_t \rightarrow G_t \rightarrow I_t$$

where G_t denotes the maritime flow network at time t and I_t denotes the rendered network image.

This stage represents an important form of feature engineering because the CNN receives an intermediate representation that explicitly incorporates geographic structure.

CNN-Based Spatial Feature Extraction

The generated flow-network images are used as inputs to a convolutional neural network. The purpose of the CNN is to learn spatial characteristics of maritime congestion, such as whether vessel activity is concentrated within a narrow region or dispersed across several shipping areas.

The proposed CNN architecture converts the flow-network images into tensors and applies convolutional layers using Conv2d, followed by pooling operations such as MaxPool2d. These operations allow the network to identify local spatial structures while reducing image dimensionality and noise. Adaptive average pooling is subsequently used to transform the learned spatial representation into a compact feature vector.

$$I_t \xrightarrow{\text{CNN}} z_t$$

where I_t is the maritime network image and z_t represents the learned spatial features. The CNN therefore does not simply estimate the number of ships; instead, it is designed to learn the morphology of maritime traffic, including the concentration, dispersion, and spatial configuration of vessel activity.

Synthetic Training Data

A significant challenge in satellite-based maritime forecasting is the limited number of historical SAR observations available for supervised machine learning. The methodology therefore introduces synthetic flow-network images generated from patterns observed in the real SAR-derived vessel distributions.

The real observations provide the empirical basis for the synthetic-data generation process. The synthetic images are intended to expand the range of congestion states available during model training while preserving the spatial structure identified from observed maritime activity.

Observed SAR $\rightarrow$ Observed Flow Networks $\rightarrow$ Synthetic Flow Networks $\rightarrow$ CNN Training

Temporal Forecasting of Maritime Congestion

While the CNN extracts spatial information from individual observations, maritime traffic is inherently dynamic. The next stage therefore incorporates a temporal model to capture transitions between congestion states.

After the CNN generates spatial features z_t , a sequence of historical maritime states can be provided to a recurrent model such as an LSTM:

$$(z_{t-k}, \dots, z_{t-1}, z_t) \xrightarrow{\text{LSTM}} \hat{z}_{t+1}$$


The objective is to estimate the next maritime congestion state using information contained in previous observations. The project proposes using past classifications of traffic tightness and spatial patterns as inputs to an LSTM in order to predict the next vessel-cluster configuration.

Overall Model Architecture

The complete research pipeline can be summarized as follows:

1. Sentinel-1 SAR imagery acquisition
2. Ship-candidate detection from radar backscatter
3. Regional vessel-count construction
4. Flow-network image generation
5. CNN-based spatial feature extraction
6. LSTM-based temporal modeling
7. Forecast of the next maritime congestion state

Sentinel-1 SAR $\rightarrow$ Ship Detection $\rightarrow$ Regional Vessel Counts $\rightarrow$ Flow-Network Images $\rightarrow$ CNN Features $\rightarrow$ LSTM $\rightarrow$ Next Maritime State



Architecture:

Next, we convert congestion into a graph/network image since we want to analyze the concentration of clusters that exist in the shipping for the time series images. And we look at the traffic structures to design graphic network and image representations of maritime flow out of it.

Next, once we have the graphic flow, then we design so that CNN learns this traffic morphology and generate synthetic training data. Furthermore, after we use the real time series satellite data and transform into flow-network images as synthetic generation grounded upon the SAR observed maritime activity. For the flow network, this is the feature engineering where we turn the raw data into a representation that CNN can easily understand. So this allows using statistics to extract signals and into complex patterns.

Since CNN now reads the render positions where each node is the locations of detected positions for distribution of ships, we can get information about how tight or wide the spread of the ships are from one another. And using torch.tensor, CNN converts spatial inputs into tensors. Then, we use nn.Conv2d into 16 small spatial filters to answer where the spatial structures exist locally. Next, we use nn.MaxPool2d(2) to shrink the image and reduce the noise level and CNN learns from nn.Conv2d(x,y,z, ...) iterations of patterns to see the locations of the spatial dispersion and use AdaptiveAvgPool to vectorize the spatial information.

Then we use argmax() to learn the feature inputs of spread. Since the congestion buildup is happening, we capture the transition dynamics and use the results of classifications about probabilities of tightness across the lanes and extract visual patterns using LSTM and feed the past maritime states to predict the next cluster of the images.

Research Contribution

The central contribution of the framework is the transformation of raw satellite imagery into a structured spatiotemporal maritime signal. Instead of treating vessel counts as a purely numerical time series, the methodology preserves geographic information describing where congestion occurs and then applies CNN-based representation learning to identify spatial patterns.

Within this framework, the CNN addresses the question, “What does the current spatial structure of maritime traffic look like?” The temporal model then addresses the complementary question, “Given recent maritime states, what is the likely next congestion state?”

The resulting framework creates a potential bridge between remote sensing, computer vision, time-series forecasting, and oil-market and logistics analysis.

Terminology Recommendation

For a formal research paper, the phrase “predict the next clusters of ships” can be replaced with “forecast the spatial distribution of maritime congestion.” This wording is more precise because the model is forecasting a future traffic configuration or congestion state rather than the behavior of individual vessels.

Source Basis

This draft is based on the user’s two project documents: (1) the technical implementation notebook containing the Sentinel-1 SAR acquisition, vessel-detection, regional-count, flow-network, and machine-learning code; and (2) the research description explaining the use of CNN and LSTM methods for maritime congestion analysis.