Proceedings of the Annual General Donald R. Keith Memorial Conference
West Point, New York, USA
April 23, 2026
A Regional Conference of the Society for Industrial and Systems Engineering

Identifying Logistics Patterns Associated with Precursor Chemical Shipments in Global Supply Chains

Gabriel Bush, Jackson Garrett, Luke Guajardo, Nicholas Thliveris, Lt Col Taylor Leonard, Gerry Gonzalez, and Benjamin Kallemyn

Operations Research Program, United States Air Force Academy, Colorado Springs, CO 80841

Corresponding author's Email: gerry.gonzalez@afacademy.af.edu

Author Note: Gabriel Bush, Jackson Garrett, Luke Guajardo, and Nicholas Thliveris are first class cadets at the U.S. Air Force Academy (USAFA) majoring in Operations Research and Lt Col Taylor Leonard, Gerry Gonzalez and Benjamin Kallemyn are faculty advisors. MAJ Verhotz, and Mr. Angel Cordero work in the J391 division of U.S North Command (USNORTHCOM) in Peterson Space Force Base located in Colorado Springs, Colorado. The views expressed herein are those of the authors and do not reflect the position of the U.S. Air Force Academy, the Department of the Air Force, or the Department of War.

Abstract: USNORTHCOM J391 Counter Threat Chemicals Operations monitors global maritime trade to identify risks from precursor chemical movement. This project developed a data-driven framework using 50,523 shipment records (2011–2025) to detect logistics patterns linked to declared precursor chemicals. A supervised Random Forest model, trained on attributes such as ports, vessels, shippers, consignees, and carriers, predicts a proxy risk label with high accuracy (mean ROC-AUC 0.969). An interactive application enables data upload, model training, and shipment scoring, providing calibrated risk scores and feature-level explanations to support analyst interpretation. While not directly detecting illicit trafficking, the system offers a scalable decision-support tool that helps decision-makers prioritize shipments for review based on logistics patterns consistent with precursor-related movement, enhancing USNORTHCOM J391's operational effectiveness.

Keywords: precursor chemicals · supply chain analytics · logistics pattern analysis · machine learning · Random Forest · USNORTHCOM · decision-support systems · maritime trade analytics

1. Introduction

North American Aerospace Defense Command (NORAD) and United States Northern Command (USNORTHCOM) operate as a co-located headquarters at Peterson Space Force Base in Colorado Springs. Within USNORTHCOM, the J39 Plans and Policy Directorate includes the J391 Counter-Transnational Criminal Organizations (CTCO) branch, which supports homeland defense by helping identify illicit networks that exploit legitimate global supply chains. A growing challenge for this mission involves precursor chemicals used in synthetic drug production. Transnational criminal organizations move these chemicals through legitimate maritime trade by routing shipments across multiple countries, embedding them within normal commercial cargo, and exploiting inconsistencies in shipping documentation. These tactics make early detection difficult and create large volumes of trade data that analysts must manually review.

This project develops a data-driven framework to identify shipments that exhibit logistics patterns associated with declared precursor-chemical activity. Rather than relying on potentially unreliable declaration data, the approach analyzes Bill of Lading records that capture observable shipping behavior such as ports of lading and unlading, vessel usage, and commercial relationships among logistics entities. Using these logistics attributes, a machine-learning model identifies shipments whose routing behavior deviates from normal trade patterns and assigns a risk score for analyst review. The model is designed as a decision-support tool that complements keyword-based approaches by leveraging logistics-network structure, rather than directly identifying illicit activity for the USNORTHCOM J391 CTCO mission.

1.1 Problem Statement

The trafficking of precursor chemicals used in synthetic drug production presents a significant challenge to U.S. homeland security due to the ability of transnational criminal organizations to exploit legitimate global supply chains. These organizations embed chemical movement within normal commercial activity by rerouting shipments through multiple ports, repeatedly using trusted transport assets, and leveraging established commercial intermediaries to obscure traditional risk indicators. As a result, potentially sensitive shipments often appear benign within large volumes of lawful trade data. While commodity declarations and shipment descriptions provide useful signals, they are susceptible to mislabeling, vague classification, and inconsistent reporting practices. Bill of Lading data offers a more robust foundation for analysis by capturing the physical movement of goods and the structure of logistics networks, including routing patterns, transport assets, and recurring commercial relationships. These observable logistics characteristics provide an alternative lens for identifying patterns associated with precursor-related activity. However, the absence of ground-truth data linking shipments to confirmed illicit activity limits the ability to directly detect trafficking behavior. Instead, this project focuses on identifying shipments whose logistics patterns resemble those associated with declared precursor-chemical activity, using a proxy label derived from keyword-based identification. This approach enables the analysis of supply-chain structure while acknowledging the limitations of available data. The problem addressed in this project is to identify shipments whose logistics patterns resemble those associated with declared precursor-chemical activity, enabling improved prioritization of analyst review within large volumes of global trade data and supporting more efficient allocation of investigative resources within the USNORTHCOM J391 CTCO mission.

1.2 Related Work

Recent research demonstrates how artificial intelligence, data analytics, and supply-chain modeling can improve the detection of illicit activity hidden within legitimate commercial systems. Hu et al. (2024) show that large language models and automated text-analysis tools can identify coded or concealed drug-trafficking conversations on social media. Their findings highlight the ability of machine learning systems to surface weak signals within large, complex datasets, illustrating how automated pattern recognition can assist analysts in identifying criminal activity that would otherwise remain hidden.

Ayaz (2024) examines the global fentanyl supply chain linking China, Mexico, and the United States and identifies structural vulnerabilities within customs enforcement, transshipment hubs, and international logistics networks. The study demonstrates how precursor chemicals can move through legitimate trade routes while supporting illicit drug production. This work highlights the importance of analyzing upstream chemical flows rather than focusing solely on finished narcotics trafficking.

Morris (2013) analyzes how transnational criminal organizations exploit weaknesses in governance, corruption, and international trade systems to conceal illicit activity within legitimate commercial networks. His work provides geopolitical context for understanding how recurring trade routes, commercial intermediaries, and logistics providers can become embedded in trafficking networks.

Together, these studies highlight two important insights relevant to this project. First, illicit activity frequently hides within large volumes of legitimate commercial data, making machine-learning approaches useful for identifying subtle behavioral patterns that may indicate risk. Second, the structure of international supply chains plays a critical role in enabling illicit chemical movement, meaning that logistics data such as ports, vessels, and commercial relationships can provide valuable indicators of suspicious activity.

These insights inform the analytical approach used in this project. Rather than relying solely on self-reported chemical declarations, the model analyzes logistics-based shipment characteristics derived from Bill of Lading records. A supervised machine-learning model evaluates patterns in shipping routes, vessel usage, and commercial relationships to identify shipments whose behavior deviates from normal trade patterns and may warrant further investigation by USNORTHCOM J391 CTCO analysts.

2. Data and Methodology

This project analyzes 50,523 shipment-level observations from Bill of Lading records collected between 2011 and 2025. Each observation represents a unique import shipment and includes logistics attributes such as foreign ports of lading, U.S. ports of unlading, vessel name, carrier code, shipper, consignee, shipment weight, and commodity description. These variables capture routing behavior and commercial relationships within global maritime supply chains.

A binary risk indicator is constructed to identify shipments associated with declared precursor chemicals used in synthetic drug production. Commodity descriptions are matched against a curated list of precursor chemical keywords derived from client guidance and open-source intelligence. Shipments containing these keywords are labeled as potential precursor shipments, resulting in approximately one percent of the dataset being classified as high risk.

This labeling approach creates a proxy outcome rather than a confirmed measure of illicit activity. The model therefore learns patterns associated with shipments that explicitly declare precursor-related content, rather than detecting concealed or misrepresented trafficking behavior. As a result, model performance should be interpreted as the ability to recover logistics patterns associated with declared precursor shipments, not as direct detection of illegal activity.

Shipment attributes used by the model include foreign port of lading, U.S. port of unlading, vessel name, carrier code, shipper, and consignee. These categorical variables are converted into numerical representations using label encoding so they can be used by the machine-learning model. While Random Forest models can handle high-cardinality categorical variables, features such as vessel, shipper, and consignee function as identifiers and may introduce risk of overfitting to specific entities. Future work may explore aggregation or embedding-based representations to improve generalizability across unseen logistics actors.

2.1 Formulation

The objective of this project is to develop a supervised learning model that predicts whether a shipment contains declared precursor-chemical indicators based on observable logistics behavior. Each shipment is represented as a feature vector $x_i$derived from Bill of Lading records, including foreign port of lading, U.S. port of unlading, vessel name, carrier code, shipper, and consignee. These variables capture patterns of routing structure, transport asset utilization, and recurring commercial relationships within global trade networks. The dependent variable is a binary label derived from commodity descriptions using keyword matching. For each shipment i, the response variable is defined as:

$$y_i = \begin{cases} 1 & \text{if the shipment contains a precursor-chemical keyword indicator} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

This formulation reflects a proxy labeling strategy rather than a direct measure of illicit activity. The model therefore estimates the conditional probability:

$$P(y_i = 1 \mid x_i) \quad (2)$$

which represents the likelihood that a shipment's logistics characteristics are consistent with those of declared precursor shipments. Given the binary structure of the response variable and the extreme class imbalance present in the dataset, a Random Forest classifier is selected. Random Forest models are well suited for this application because they capture nonlinear relationships, handle high-dimensional categorical inputs, and maintain stable performance under imbalanced conditions. The Random Forest consists of Mdecision trees, each trained on a bootstrap sample of the dataset. For a shipment with feature vector x_i, each tree $T_m$produces an estimate of $P(y_i = 1 \mid x_i)$. The ensemble prediction is defined as:

$$p_i = \frac{1}{M}\sum_{m=1}^{M} T_m\,(x_i) \quad (3)$$

This probability is converted into an operational risk score:

$$RiskScore_i = 100 \times p_i \quad (4)$$

which allows analysts to interpret model output as a percentage-based indicator of similarity to declared precursor shipment patterns. These scores support threshold-based prioritization for analyst review rather than automated enforcement decisions.

2.2 Methodology

Although precursor-chemical trafficking can be framed as an anomaly-detection problem, the availability of a curated keyword list enables a supervised learning approach focused on identifying patterns associated with declared precursor shipments. Rather than detecting purely statistical outliers, the model learns logistics characteristics that distinguish these shipments from general trade activity.

Data preparation transforms Bill of Lading records into a machine-learning–ready format while preserving their operational meaning. Logistics-based categorical variables, including ports, vessels, and commercial entities, are encoded numerically using label encoding. This approach allows the Random Forest model to incorporate complex relationships across supply-chain actors while maintaining interpretability. The model is evaluated using stratified cross-validation to preserve the approximately one-percent prevalence of high-risk observations across training and validation folds. Performance is assessed primarily using ROC-AUC, with additional consideration of precision, recall, and F1-score due to the highly imbalanced nature of the dataset.

For each shipment, the model generates a predicted probability that the observation belongs to the high-risk class. This probability is converted into a percentage-based risk score that supports operational screening. Shipments with elevated risk scores can be prioritized for analyst review.

The model is designed as a decision-support tool rather than a standalone detection system. By identifying shipments whose logistics patterns resemble those of declared precursor shipments, the framework complements keyword-based approaches and enables analysts to focus attention on a smaller subset of potentially relevant records within large trade datasets.

ISBN: 97819384962-5-7

2.3 Implementation and Advantages

The shipment risk-scoring system is implemented using Python to ensure transparency, reproducibility, and ease of integration with existing analytical workflows. The analytical pipeline is built around a Random Forest classifier trained on logistics-based shipment attributes derived from Bill of Lading records. These attributes include foreign port of lading, U.S. port of unlading, vessel name, carrier code, shipper, and consignee.

After the dataset is uploaded and preprocessed, the system trains the Random Forest model using stratified cross-validation to account for the extreme class imbalance present in the data. The trained model then generates probability estimates for each shipment, representing the likelihood that the shipment contains characteristics consistent with precursor-chemical trafficking behavior. These probabilities are converted into percentage-based risk scores that allow analysts to quickly interpret model outputs and apply threshold-based screening rules.

To make the system operationally accessible, the analytical framework is integrated into an interactive application developed using the Stream lit web framework. The interface allows analysts to upload historical shipment data, train the model, score individual shipment observations, and perform batch scoring on larger datasets. The application also displays model performance metrics and feature-importance outputs that help analysts interpret the results.

The resulting system provides several operational advantages. First, the machine-learning model enables large shipment datasets to be screened automatically, dramatically reducing the volume of records requiring manual review. Second, the use of logistics-based variables allows the system to identify logistics patterns associated with declared precursor shipments even when commodity descriptions are incomplete or inconsistent. Third, the interactive interface allows analysts to apply the model without requiring advanced programming knowledge. Together, these capabilities create a scalable decision-support tool that enhances early detection of precursor-chemical trafficking activity within global maritime trade networks.

2.4 Limitations

This labeling approach introduces several limitations that warrant consideration. The identification of precursor-related shipments relies on matching known chemical names within free-text commodity descriptions. As a result, shipments that intentionally obscure chemical identity through vague terminology, abbreviations, or alternative naming conventions may not be detected, potentially leading to false negatives. Conversely, legitimate industrial shipments that contain lawful precursor chemicals may be flagged despite benign intent, resulting in false positives.

Another limitation arises from the construction of the training labels themselves. The labeling process captures only shipments that contain known precursor indicators present in the curated chemical keyword list. Consequently, the dataset represents only a partial view of potential precursor-chemical trafficking activity. Shipments involving previously unidentified chemicals or alternative concealment strategies may not be represented within the labeled data.

The dataset also exhibits extreme class imbalance, with approximately one percent of observations labeled as high risk. While stratified evaluation techniques help mitigate this issue, the imbalance reinforces that the model learns from a limited subset of potentially suspicious shipments.

For these reasons, the model should be interpreted as a prioritization and early-warning tool rather than a definitive interdiction mechanism. The system is designed to assist analysts by highlighting shipments that exhibit logistics patterns consistent with precursor-related activity, enabling more efficient allocation of investigative resources within the USNORTHCOM J391 CTCO mission.

3. Results

Model performance is evaluated using stratified cross-validation to ensure stability under extreme class imbalance. The Random Forest classifier achieves a mean ROC-AUC of 0.969 with a standard deviation of 0.012, indicating strong discriminative performance and consistent generalization across validation folds. These results demonstrate that the model effectively distinguishes shipments containing declared precursor chemical indicators from general trade activity using logistics characteristics derived from Bill of Lading records. However, this performance reflects the model's ability to predict the proxy label derived from keyword-based identification, rather than confirmed illicit trafficking activity. As a result, the evaluation should be interpreted as measuring how well logistics patterns explain declared precursor shipments, not direct detection of illegal behavior. Additional evaluation metrics including precision, recall, and F1-score are considered to account for the highly imbalanced dataset.

When applied to the full Bill of Lading dataset, the model generates a probability estimate for each shipment representing the likelihood that the observation exhibits logistics patterns consistent with declared precursor shipments. These probabilities are converted into percentage-based risk scores to support operational interpretation and prioritization. Using a 50 percent threshold, the model flags 484 shipments out of 50,523 total records, corresponding to approximately 0.96 percent of all shipments. This result aligns with the prevalence of the proxy label and reflects the conservative nature of the classification threshold.

Analysis of model behavior indicates that classification is driven primarily by logistics-network structure rather than commodity descriptions alone. The most influential predictors include foreign port of lading, U.S. port of unlading, vessel name, shipper, consignee, and carrier code. These variables capture recurring routing patterns, commercial relationships, and transport asset usage across global trade networks. This finding supports the use of Bill of Lading data as a meaningful source of signal for identifying patterns associated with declared precursor shipments.

To translate model outputs into operational insight, high-risk shipments were aggregated by recurring routing characteristics. Table 1 lists the most frequent foreign ports of lading among shipments flagged at or above the 50 percent threshold. The concentration of risk within a small number of origin ports suggests that upstream monitoring and interdiction efforts may be prioritized along specific international trade corridors.

Table 1. Most Frequent Foreign Ports of Lading Among High-Risk Shipments (Risk ≥ 50%)

Rank	**Foreign Port of Lading**
1	Shanghai, China (Mainland)
2	Yangshan, China (Mainland)
3	Veracruz, Mexico
4	Algeciras, Spain
5	Pusan, South Korea

A similar pattern appears on the domestic side of the supply chain. Table 2 presents the most frequent U.S. ports of unlading among shipments flagged at or above the 50 percent threshold. These ports represent key domestic chokepoints where inspection resources, interagency coordination, and targeted screening efforts may yield the greatest operational benefit.

Table 2. Most Frequent U.S. Ports of Unlading Among High-Risk Shipments (Risk ≥ 50%)

Rank	**U.S. Port of Unlading**
1	Newark, New Jersey
2	Houston, Texas
3	New York, New York
4	Tacoma, Washington
5	Seattle, Washington

These geographic concentrations illustrate how machine learning can highlight systemic trade patterns rather than only isolated suspicious shipments. By identifying clusters of elevated risk within specific shipping corridors, the framework provides analysts with additional context that can inform monitoring priorities and investigative focus.

To further contextualize the operational importance of precursor-chemical detection, Table 3 summarizes recent drug overdose statistics reported by the Centers for Disease Control and Prevention (CDC) and the National Center for Health Statistics (NCHS). These figures illustrate the scale of the synthetic-drug crisis in the United States and highlight the importance of upstream detection efforts aimed at disrupting precursor supply chains before narcotics reach downstream distribution networks.

Table 3. U.S. Drug Overdose Deaths (CDC/NCHS Estimates)

Year	**All Drug Overdose Deaths**	**Opioid-Involved Overdose Deaths**
2023	110,037	83,140
2024	80,391	54,743

While the model does not directly measure downstream outcomes such as overdose rates or confirmed trafficking activity, these statistics provide context for the broader operational importance of monitoring precursor chemical movement. The model focuses specifically on identifying logistics patterns associated with declared precursor shipments, which may support upstream analytical prioritization within larger counter-drug efforts.

Overall, the results demonstrate that the model identifies a small subset of shipments whose logistics patterns closely resemble those of declared precursor shipments. By reducing a dataset of more than fifty thousand records to a manageable subset requiring analyst attention, the framework supports more efficient prioritization and reduces manual screening burden. The model should be interpreted as a tool for highlighting structured patterns in logistics data and supporting analyst decision-making, rather than a system for directly identifying illicit trafficking activity.

4. Conclusion and Future Research

This project develops and evaluates a machine-learning–based risk-scoring framework that identifies logistics patterns associated with shipments containing declared precursor-chemical indicators. By leveraging Bill of Lading data, the approach demonstrates that supply-chain structure contains predictive signal beyond commodity descriptions. The model achieves strong performance in predicting this proxy label and provides a practical mechanism for prioritizing analyst attention within large volumes of global trade data. The model does not directly detect illicit trafficking activity, as no ground-truth dataset linking shipments to confirmed criminal behavior is available. Instead, it provides an intermediate analytical capability that highlights logistics patterns consistent with declared precursor shipments. This distinction is important, as model outputs should be interpreted as indicators for further investigation rather than definitive evidence of illicit activity.

To support operational use, the framework is implemented in an interactive application that enables analysts to upload data, train the model, and generate shipment-level risk scores. This design improves accessibility, supports integration into existing workflows, and allows both individual and batch analysis of shipment data.

Overall, this work provides a scalable and operationally relevant decision-support capability for USNORTHCOM J391 CTCO. The framework enhances consistency in shipment prioritization, reduces manual screening burden, and improves the ability to focus analytical resources on shipments whose logistics patterns warrant further review. Future research may incorporate additional data sources, including intelligence-based labels, and explore alternative feature representations to improve generalizability and better support identification of concealed or misrepresented activity.

5. References

Ayaz, O. (2024). *Global fentanyl supply chains: Chemical sourcing, trafficking routes, and enforcement vulnerabilities.* Journal of Transnational Crime Studies, 18(2), 45–63.

Hu, X., Zhang, Y., Reid, A., & Lopez, M. (2024). Detecting concealed narcotics-trafficking conversations using large language models. *International Journal of Illicit Networks and Analytics, 12*(1), 55–78.

Morris, S. (2013). *Cartel dynamics and governance vulnerabilities in North American trade systems.* Cambridge University Press.

OpenAI. (2025). *ChatGPT (Version 5.1)* [Large language model]. https://www.openai.com/