big data and cognitive computing

MDPI

Article

Harnessing Graph Neural Networks to Predict International Trade Flows

Bassem Sellami [1,*], Chahinez Ounoughi [2], Tarmo Kalvet [2], Marek Tiits [2] and Diego Rincon-Yanez [3,4]

[1] Department of Software Science, Tallinn University of Technology, Akadeemia tee 15a, 12618 Tallinn, Estonia
[2] Department of Business Administration, Tallinn University of Technology, Ehitajate tee 5, 12618 Tallinn, Estonia; chahinez.ounoughi@taltech.ee (C.O.); tarmo.kalvet@taltech.ee (T.K.); marek.tiits@taltech.ee (M.T.)
[3] Department of Information and Electrical Engineering and Applied Mathematics, University of Salerno, Via Giovanni Paolo II 132, 84084 Fisciano, Italy; drinconyanez@unisa.it
[4] Facultad de Ingenierías y Tecnologías, Universidad de Santander, Cúcuta 540001, Colombia
* Correspondence: bassem.sellami@taltech.ee

Abstract: In the realm of international trade and economic development, the prediction of trade flows between countries is crucial for identifying export opportunities. Commonly used log-linear regression models are constrained due to difficulties when dealing with extensive, high-cardinality datasets, and the utilization of machine learning techniques in predictions offers new possibilities. We examine the predictive power of Graph Neural Networks (GNNs) in estimating the value of bilateral trade between countries. We work with detailed UN Comtrade data that represent annual bilateral trade in goods between any two countries in the world and more than 5000 product groups. We explore two different types of GNNs, namely Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs), by applying them to trade flow data. This study evaluates the effectiveness of GNNs relative to traditional machine learning techniques such as random forest and examines the possible effects of data drift on their performance. Our findings reveal the superior predictive capability of GNNs, suggesting their effectiveness in modeling complex trade relationships. The research presented in this work offers a data-driven foundation for decision-making and is relevant for business strategies and policymaking as it helps in identifying markets, products, and sectors with significant development potential.

Keywords: international trade; machine learning; random forest; graph neural networks; graph convolutional networks; graph attention networks; data drift

check for updates

Citation: Sellami, B.; Ounoughi, C.; Kalvet, T.; Tiits, M.; Rincon-Yanez, D. Harnessing Graph Neural Networks to Predict International Trade Flows. *Big Data Cogn. Comput.* **2024**, *8*, 65. https://doi.org/10.3390/bdcc8060065

Academic Editor: Fabrizio Marozzo

Received: 22 April 2024
Revised: 22 May 2024
Accepted: 3 June 2024
Published: 7 June 2024

Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The gravity model of international trade [1,2] is a common empirical method for predicting bilateral trade flows between nations. Gravity models rely on key economic metrics such as gross domestic product and the distance between countries of origin and destination, indicating that larger countries are more likely to participate in trade, while trade flows tend to decline as the distance between countries increases. Additionally, in earlier works, additional factors such as the existence of common borders, cultural proximity, or the presence of trade agreements have been considered to analyze trade between countries. Generally, gravity models have been estimated using conventional statistical techniques like log-linear regressions. Yet, these methods encounter difficulties when dealing with extensive, high-cardinality datasets, exemplified by the thousands of product codes in international trade dataset.

The utilization of machine learning techniques in estimating gravity models of international trade represents a significant shift, offering new possibilities. Recent research [3] has shown that a machine learning method based on random forest (RF) can accommodate thousands of product groups within a single integrated model, thereby predicting bilateral

trade flows with high precision at the detailed six-digit HS product code level. However, alongside the RF methodology, other approaches are rapidly gaining traction and proving to be valuable in the fields of data science and predictive analytics. Specifically, Graph Neural Networks (GNNs) are emerging as a standout technology [4,5].

GNNs are uniquely adept at handling complex, irregular data structures found in graphs. This makes GNNs particularly powerful for applications involving relational data, such as global economic and social networks, where the connections between entities can provide critical insights for prediction tasks. GNNs also stand out for their ability to model dynamic networks [6].

The hypothesis of this research article is that GNNs offer superior predictive capabilities for modeling international trade flows over traditional machine learning models like Random Forests due to their inherent ability to capture the complex, non-linear relationships and structural dependencies present within global trade relations.

Our research methodically assesses the predictive potential of these models across diverse scenarios. We employ a variety of input feature combinations encompassing geographical, economic, and temporal variables to ensure a comprehensive evaluation and offer a complete view of their performance. Furthermore, our study introduces novel assessment dimensions and extends the gravity model application to a global scale, surpassing previous research.

This study employs empirical analysis to validate the efficiency and practicality of the proposed approach. It is aimed at delivering accurate predictions for bilateral trade figures at the level of individual product groups and countries, thus providing policymakers, businesses, and other relevant stakeholders with essential insights. These insights are crucial for refining trade policies, uncovering novel market opportunities, stimulating economic growth, and adapting to the dynamic landscape of international trade. The ability to accurately forecast bilateral export opportunities is critically important for entrepreneurs and policymakers in small and medium-sized economies who aim to achieve sustainable economic development [7,8].

Thus, the research plays a vital role in shaping business strategies and policymaking by identifying markets, products, and sectors with significant potential. Recognizing such areas can lead to a more effective allocation of resources to enhance infrastructure, finance mechanisms, and knowledge development. Furthermore, the application of machine learning models to predict trade flows and market dynamics offers a data-driven foundation for decision-making, reducing the uncertainty associated with management and governance.

The paper is organized as follows: Section 2 reviews the relevant literature, focusing on the enhancement of gravity models of international trade through random forest techniques as well as prior research on the utilization of GNNs in the analysis of international trade. Section 3 details our approach for predicting trade values, employing a comprehensive dataset that merges international trade data with gravity data. Section 4 presents the results of our study. In Section 5, we evaluate the performance of GNNs, explore the impact of data drift, and assess the significance of our results for business and policy stakeholders. Section 6 summarizes the paper and suggests directions for future research.

2. Related Work

Predicting trade flows has long been a challenge in international economics that has attracted the interest of scholars worldwide. Scholars have approached this challenging subject using a variety of approaches and methodologies, ranging from conventional statistical models to modern machine-learning techniques. This section provides a synopsis of the relevant literature concerning trade flow prediction, with an emphasis on Random Forest and Graph Neural Network applications.

2.1. Gravity Model of International Trade

The gravity model of trade, a 'work horse' of international economics, was introduced by Tinbergen in 1962 [2]. It suggests that trade flows decrease as the geographic distance

between trading partners increases while increasing in tandem with their Gross Domestic Product (GDP). The basic model for trade between two countries (i and j) takes the form of

$$F_{ij} = G \cdot \frac{M_i M_j}{D_{ij}} \quad (1)$$

In this formula, F stands for trade flow, M stands for the size of the economies under consideration, D stands for the distance between countries, and G is a constant.

Since its inception, the gravity model's popularity and application in both academic and policy-making spheres have markedly grown [9,10], establishing it as a foundational element of empirical trade research. The model's robustness lies in its strong empirical predictive capabilities and simplicity. Furthermore, the model's flexibility allows for the integration of additional variables, such as shared borders [11], common languages, historical colonial relationships, and trade agreements [12], as well as factors affecting ease of conducting business and institutional frameworks [13]. This versatility facilitates comprehensive analysis across diverse settings. The gravity model is instrumental in policy evaluation, offering policymakers a tool to assess the implications of trade policy modifications on trade flows, including tariff adjustments or the introduction of free trade agreements [12,14,15]. Such a model has the capacity to provide essential insights for informed policy formulation.

The widely used Ordinary Least Squares (OLS) and Poisson regression methods provide a straightforward framework for understanding the determinants of trade. When estimating the gravity model of trade using OLS, the model is often specified in a log-linear form. This transformation has several benefits: it linearizes the relationship, it stabilizes the variance (homoscedasticity), and the coefficients can be interpreted as elasticities. One downside is that logarithms cannot be applied to zero or negative values. This is problematic in datasets with zero trade flows. Various methods, like adding a small constant to zero values or using alternative models, are employed to address this issue. When using Poisson regression to estimate the gravity model of trade, there is no need to log-transform the dependent variable; it models the trade volumes directly and can naturally handle zero trade flows, making it a robust choice for trade data.

However, conventional statistical methods are not equipped to handle large, high-cardinality datasets, like the thousands of product codes in international trade datasets. Utilizing these models through conventional approaches requires creating an extensive array of dummy features. An alternative strategy involves estimating gravity models of international trade on a per-product-group basis. This necessitates assessing, validating, and maintaining thousands of regression models to encompass the full breadth of the product spectrum.

Thus far, aggregated trade or trade in specific products has been mostly estimated with linear models, and multi-variate linear gravity models of trade are rare. In a particularly comprehensive study, 15 years of bilateral trade data, broken down into over 1200 products, were employed to develop various gravity models. The findings indicated that countries exporting products related to those of a destination, countries exporting a given product to the neighbors of a destination, or those whose neighbors are already exporting the same product tend to see an increase in their exports of that product to the destination [16].

2.2. Random Forest-Based Machine Learning Models of International Trade

In contrast, machine learning models are inherently capable of handling voluminous datasets and impose fewer assumptions about the properties of data compared to econometric models, which typically require specific assumptions. Machine learning approaches can efficiently process high-dimensional data. They exhibit greater resilience to assumption violations and possess the flexibility to adjust to changing data patterns [17,18]. Moreover, machine learning techniques offer the advantage of automated feature selection and hyperparameter tuning, thereby saving time and improving model performance.

Recently, there has been a substantial increase in the application of machine learning within economics and econometrics. Decision trees and Random Forests have demonstrated remarkable effectiveness in simulating non-linear linkages in data since Ho introduced them in 1995 [19]. A Random Forest model is a powerful ensemble learning method consisting of multiple decision trees that work together to make predictions. In a Random Forest, the predictions from multiple decision trees are combined, increasing the precision of predictions. The approach involves creating bootstrapped training sets using the original dataset and making random selections of features for each tree to encourage variance among the trees. Once each individual tree prediction is obtained, the final forecast is calculated, typically by averaging or employing majority voting in the case of binary classification.

Random forest techniques require minimal data pre-processing and are capable of handling both categorical and continuous data, making them versatile for various types of datasets. Additionally, they can accommodate non-normal data distributions, such as the Poisson distribution. One of the most recent contributions proposes a novel practical approach to employing a Random Forest for estimating a multi-product-group gravity model of international trade [3]. This model covers monthly bilateral trade between any two countries in the world and allows for the prediction of exports across more than 5000 product groups.

However, while Random Forest models are well suited for analyzing trade data, they have limitations in extrapolation. This inability to extrapolate effectively can hinder their accuracy in predicting future trade flows over time. Furthermore, in general, deep learning regression models tend to outperform Random Forest models.

2.3. Graph Neural Networks

GNNs constitute a family of techniques designed to broaden the applicability of neural networks, originally conceived for Euclidean data, to graph-structured data [20]. GNNs demonstrate their versatility in addressing diverse tasks [21], encompassing node-level operations [22], graph-level tasks, edge classification, predictive link analysis, as well as node clustering and community detection. GNNs are composed of layers that process data organized within graph structures. These models have exhibited their effectiveness across a range of applications, including document labeling, traffic forecasting, and fraud detection [23]. Recently, some research has also used GNN to address the challenges related to international trade in products and services [24–26].

The application of GNNs in the analysis of international trade involves constructing a graph where nodes represent individual countries and edges denote the relationships between them. Each node encapsulates the unique economic features and properties of a specific nation, such as gross domestic product or population. The edges represent interdependencies between nations, including trade flows, geographic distance, or terms of trade agreements. Learning the representation of nodes in the graph, referred to as node embedding, is a crucial component for downstream tasks such as classification and regression. The majority of node embedding models are built upon either spectral decomposition techniques [22,27] or matrix factorization methods [28,29]. However, it should be noted that most embedding frameworks are inherently transductive in nature and are capable of generating embeddings for a single fixed graph only. These transductive approaches lack efficiency in generalizing to unseen nodes, particularly in the case of evolving graphs. Furthermore, these approaches are unable to acquire the ability to generalize across different graphs.

New methods utilizing causality have been recently introduced to overcome these limitations and enhance the model's generalization capabilities. For example, Minakawa et al. [25] proposed a GGAE (Gravity-informed Graph Auto-encoder) and its surrogate model, which draws inspiration from the gravity model. They demonstrated that the surrogate model can solve the edge weight prediction problem in GNNs and that the gravity model's prediction of trade value can be expressed as an edge weight prediction problem [30]. Moreover, Monken et al. [26] introduced AINET, a method based on GNN that is tailored for bilateral

trade modeling and used to measure causal scenarios during outlier events in international trade, providing explainability and predicting outliers effectively.

Likewise, Ahmed et al. [31] introduced a GNN-based method for modeling product relationships and predicting unseen product networks. Products are represented as nodes and relationships as edges in a graph, using GraphSAGE to learn continuous node and edge representations. This approach integrates product features with a classification model for predicting product relationships.

On the other hand, Panford-Quainoo et al. [32] examined the relationship between the trade gravity model and GNNs. They framed the task of classifying a country's GDP as a node classification problem in GNNs and the task of finding trade partners as a link prediction problem in GNNs. However, they did not address the prediction of trade value, which can be viewed as an edge weight prediction problem in GNNs. Verstyuk et al. [24] argued that GNNs are a natural and theoretically appealing class of models for international trade. For this claim, they provide evidence for it by showing the theoretical link and real-world results from a large set of yearly country-level data that are used to look at two-way accessibility.

Table 1 provides an overview of various research works on various analytical and GNN-based approaches used in international trade research. It describes the methodology or model employed, the data used for analysis, and the specific or defined focus of each study. The research objectives outlined range from statistical classification and trade forecasting to analysis of trade effects and forecasting communication loads and trade flow. This collection of works highlights the evolving nature of trade analytics, incorporating traditional economic modeling with state-of-the-art machine learning techniques to better understand and predict patterns in international trade.

Table 1. GNN-based approaches.

Article	Approach/Model	Dataset	Specification
Monken et al. [26]	Artificial Intelligence Network Explanation of Trade (AINET) based GNN	UN Comtrade	Measure causal scenarios during outlier events in international trade.
Verstyuk et al. [24]	Analyze international trade relationships based on the gravity model using GNN model	CEPII Gravity	Create flexible versions of the traditional gravity equation and develop interpretable models
Rincon-Yanez et al. [30]	Knowledge Graph, GNN, Random Forest	CEPII Gravity	Edge weight prediction
Minakawa et al. [25]	GGAE (Gravity-informed Graph Autoencoder)	Data from the Economic Community of West African States (ECOWAS)	Trade amount prediction
Panford-Quainoo et al. [32]	Graph Convolutional Network (GCN), Graph Attention Network (GAT), Graph autoencoder (GAE) and Variational Graph Autoencoder (VGAE)	UN Comtrade	Link prediction and classification of countries
Ahmed et al. [31]	Modeling relationships between products using GNN and making predictions for unseen product networks.	Data provided by the Ford Motor Company	Link prediction of products

Building on the foundation laid by the array of approaches summarized in Table 1, the next section of our study introduces an innovative modeling strategy designed to further enhance predictive precision in international trade analytics. Specifically, Section 3 will elaborate on our adoption of GCN and GAT models. These advanced graph neural network models leverage the complex networked nature of trade data, aiming to more accurately forecast trade values by capturing and analyzing the intricate relationships and dependencies inherent in international trade structures. This methodological shift marks a

significant step forward from the traditional and current GNN-based methods discussed previously, paving the way for a more detailed discussion of our data collection and feature engineering processes.

3. Methodology

In the initial phase of this section, we provide a detailed account of the data collection process and outline the sources. Subsequently, we delve into the feature engineering procedures, detailing the techniques employed to transform raw data into informative features conducive to model training and analysis.

3.1. Data Collection and Preprocessing

The efficacy and accuracy of the trade flow forecast significantly depend on the quality of the dataset and the preparation methods used for analysis. In the following, we elaborate on the data sources, which include the UN Comtrade and CEPII Gravity datasets, as well as the pretreatment procedures involved.

3.1.1. Data Source

The dataset utilized in this study is an effective fusion of the UN Comtrade database (United Nations Comtrade. 2023. https://comtradeplus.un.org/ accessed on 6 May 2024) with the CEPII gravity data (CEPII Gravity Database, Version: CEPII,2022-2, http://www.cepii.fr/ accessed on 7 February 2023), and the World Bank data (World Bank Portal, https://databank.worldbank.org/ accessed on 14 December 2023), presenting a comprehensive and diverse information source for predicting trade flows.

- The United Nations Comtrade Database includes comprehensive records of global trade flows, including import and export information for various goods and commodities. This dataset provides detailed information on the kinds of products and services that are transferred between nations. It contains information on the trading nations engaged and product codes, quantities, and values.
- CEPII Gravity Database provides the data required to estimate the influence of geographical distance on international trade between nations.
- World Bank Database provides a comprehensive array of global economic and social statistics. This includes vital indicators such as the World Development Indicators, International Debt Statistics, and Millennium Development Indicators. Moreover, it provides extensive data on key areas such as poverty, education, and gender.

The depth and scope of the UN Comtrade dataset are improved by joining it with the gravity and World Bank features. The resulting hybrid dataset, which includes information on traded products as well as macroeconomic and geographic variables, provides a basis for a thorough understanding of global trade flows. It is the basis of our analysis for trade flow prediction.

3.1.2. Preprocessing Steps

In the following, we describe the specific preprocessing steps that were taken when combining the UN Comtrade and gravity datasets, as shown in Figure 1.

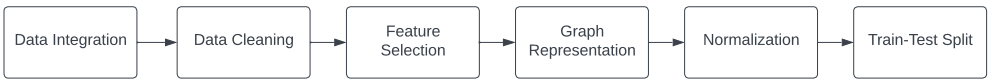


Figure 1. Preprocessing steps.

(a) Data Integration: A particular alignment is essential, ensuring the harmonization of diverse datasets based on shared factors such as commodity and country codes. This alignment is a critical prerequisite for successfully merging the Gravity features and the UN Comtrade data. By aligning these datasets, we create a unified and cohesive framework that enables the integration of information about traded goods and their

specific characteristics with broader gravity model features. This integrated dataset forms the foundation for our analysis, facilitating a comprehensive examination of the interplay between macroeconomic indicators and commodity-specific details in international trade dynamics.

The resulting hybrid dataset includes a range of factors, such as:

- Trade value: Based on data from the UN Comtrade dataset, this term refers to the real value of products and services exchanged between two nations. It is expressed in monetary terms, US Dollars (USD).
- Product code: Harmonized System (HS) codes representing particular traded goods in the UN Comtrade database enable in-depth examination of trade connections.
- Gross Domestic Product: Each trading partner's GDP is utilized in the gravity model to represent the size of each nation's economy.
- Population: Represents the total number of residents in each trading nation. This demographic factor is essential for gauging market size and potential demand for imports.
- GDP per capita: Derived by dividing the Gross Domestic Product (GDP) by the total population, GDP per capita measures a country's economic performance per individual. It offers insights into the average economic well-being of citizens in each country and their potential purchasing power.
- Distance: represents the geographical distance between the countries.

(b) Data Cleaning: In the data cleaning process, various criteria were applied to ensure the accuracy and relevance of the dataset. Focusing on the UN Comtrade annual data for the period 2017 to 2022 under the HS 2017 revision, we carefully curated a subset, retaining export data (flowCode X) related to all customs procedure codes (CustomsCode C00), all modes of transport (MotCode 0), and all last known destinations (Partner2Code 0). Augmenting this dataset, we incorporated pertinent features from the World Bank dataset, including factors like the Gross Domestic Product (gdp) of the countries of origin and destination, the GDP per capita (gdpcap) of the exporting and importing countries, and the population of the countries of origin and destination (pop). To enhance comparability, we standardized dollar values by dividing them by 1,000,000. Moreover, we incorporate the geographical distance (dist) between the countries sourced from the CEPII gravity dataset. Table 2 displays the attributes of the ultimate dataset after cleansing, offering a thorough summary of the features and their corresponding measurements.

Table 2. Bilateral Trade Datasets.

Parameter	Exports
Period	2017–2022
Reporter Code	162
Partner Code	203
Commodity Code	5385
Size	44,175,189
Features	Trade Value, Year, Commodity code, GDP, GDP per capita, Population, Distance

(c) Feature Selection: We started by examining the correlation between trade value and various features to identify promising elements for model development. Typically, features with strong absolute correlations are pivotal in model creation. However, in this scenario, correlation values with the trade value ('primaryValue') are mainly negligible, indicating a lack of linear solid relationships with individual features. Notably, (pop_o) and (gdp_o) exhibit the strongest relative correlation with (primaryValue). Considering the multicollinearity between (gdp_o) and (pop_o), choosing one is advisable to avoid issues and enhance model interpretability. Further insights

reveal a weak positive correlation between (distw_harmonic) and (gdp_o), suggesting that larger economies are more inclined to export to distant destinations. In summary, the correlation analysis emphasizes a robust linear relationship between countries' population and GDP. However, correlations with other features are generally modest or negligible. Notably, the product of population and GDP per capita emerges as a valuable factor, not only as a substitute but also as an additional insight for refining predictive models.

(d) Graph Representation: Our methodology employs a graph representation to model the relationships within the international trade network. In this graph, trade connections are depicted as edges, while countries are represented as nodes, adhering to the structure present in the original gravity dataset. This approach preserves the foundational framework of the gravity model, emphasizing the interconnectedness of countries through trade relationships [30]. A key enhancement in our graph representation lies in the incorporation of detailed information about the specific products traded between countries. This addition introduces a nuanced layer to the graph topology, providing a more comprehensive view of the trade dynamics. Each edge signifies a trade connection and encapsulates the intricacies of the traded products, allowing for a more detailed and granular representation.

(e) Data Normalization: This critical step is indispensable in our methodology because it ensures that the numerical features are uniformly scaled across various variables, encompassing macroeconomic and product-specific factors. It plays a vital role in creating a standardized and consistent framework for analysis and mitigating potential biases arising from differences in the scales of different features. Specifically, the application of MinMax scaling to the (primaryvalue), (distance), and (gdp) columns is imperative in crafting this standardized framework, further safeguarding against potential biases arising from differences in the scales of these specific features.

(f) Train–Test Split: We employed two distinct approaches for the train–test split in our model development to comprehensively assess and validate our models' predictive performance. In the first approach, Standard 20/80% Split, we divided the dataset into a training set and a test set with a ratio of 20% for testing and 80% for training. This widely used split ratio allows a substantial amount of data to be dedicated to training the model, ensuring sufficient exposure to patterns and trends within the dataset. The second approach involved a temporal split, where data from 2017 to 2021 were utilized for training the model and data from 2022 were reserved for testing. This approach is particularly useful when there is a temporal component to the data, allowing the model to be evaluated on its ability to predict future instances based on historical patterns. The temporal split offers insights into how well the model performs on more recent data, capturing potential changes or trends that may have emerged in the later period.

3.2. Graph Construction

Recently, graphs have proven to be a powerful ally in providing meaningful insights by representing complex relations within the data. These graph-oriented approaches can be applied to various problems, specifically in the context of trade flow prediction. In this approach, the graphs can map many features, representing even complex and numerous relationships [33]. In this form and to properly harness the capabilities of Graph Attention Networks (GATs) and exploit the CEPII and the UN Comtrade database, we implement a graph construction process starting with the traditional homogeneous graph expression [34] $G = (V, E)$, where (V) stands for the countries as nodes and (E) represents trade relationships between the countries as shown in Figure 2.



Figure 2. Yearly export representation description between countries represented in a graph with the node features and edge weights.

Once a generic trade relation is captured from the datasets, extracting the node (F) and edge (W) features is the next step. These two sets of values by each value and trade relation can be seen in Figure 2 and denoted as $G = (|V|F, |E|W)$, where

- V—Set of nodes representing the countries.
- F—Set of node features, where F_i is the individual features of each country.
- E—Set of edges representing trade relationships by year.
- W—Set of edge weights, where W_{ij} is the individual weight for each relationship between each i and j country.

A complete trade operation is shown in Figure 2, with the node feature representation and the direct edge weights between the countries. The node features are country-specific economic attributes such as population, GDP, and GDP per capita; however, features like the distance between the two countries and other relevant trade-specific and geographic variables have been chosen for the edge weights. Considering the dataset sample frequency in the CEPII (yearly) and UN Comtrade (yearly) databases, the graph construction process takes each frequency and stacks the information to finally merge them into a final shape that will contain the trade operations across time and commodity codes, as shown in Figure 3.

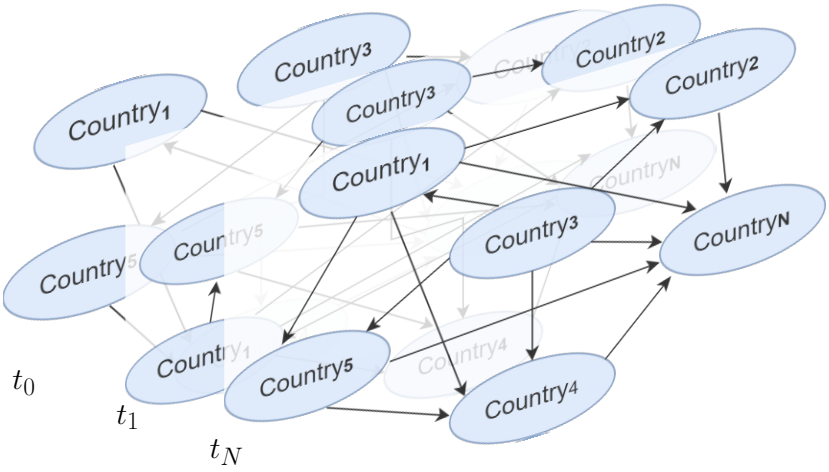


Figure 3. Isometric view of each time frame for exports in a graph shape starting from a t_0 and ending with a t_N.

By building the graph and using weights to represent the features, the ultimate goal is to turn the more than 300 features into tensors that will let the GCN and the GAT define the basic models for trade relationships.

We employed Graph Attention Networks (GAT) and Graph Convolutional Networks (GCN) to navigate and decipher the complexities inherent in the datasets. The preliminary phase involves a preprocessing where we distill over 5000 commodity codes into a finely structured format, primed for sophisticated graph-based analyses. This enabled the construction of period-specific graphs that accurately encapsulate global trade dynamics, with countries represented as nodes and their trade interactions as weighted edges. These weights ingeniously encapsulate trade volumes, geographic distances, and nuanced details of commodity exchanges, offering a granular view of international trade flows.

Utilizing the dynamic capabilities of the DGL library, we harnessed the GAT and GCN models to conduct an in-depth exploration of these graph-structured data. The GAT model, with its nuanced attention mechanism, adeptly identifies and emphasizes the most significant trade connections, allowing for a targeted analysis of pivotal trade flows. Conversely, the GCN model excels in capturing the global structure of trade networks, providing insights into the overarching patterns that govern international trade interactions.

Our methodology extends beyond mere analysis, incorporating an elaborate training process designed to capture the evolving nature of trade relationships over time. This not only ensures our models are attuned to the temporal shifts in trade dynamics but also enhances their predictive capabilities. By balancing efficiency and scalability, we manage the computational demands of processing high-dimensional data, ensuring that our approach remains both robust and practical.

3.3. Model Training

In the context of our study on predicting international trade flows using Graph Neural Networks (GNNs), selecting Random Forest as a benchmark for comparison is warranted based on its established superiority over conventional gravity models. Through this comparison, we seek to evaluate how effectively GNNs leverage their advanced capabilities in capturing complex relational structures to enhance predictive accuracy within the realm of international trade analysis.

- Random Forest (RF): A RF-based study [3] provided compelling evidence supporting the efficacy of Random Forest models over traditional gravity models in trade prediction tasks. The authors found that Random Forest models demonstrate superiority in capturing the non-linear dynamics of international trade relationships. Therefore, by leveraging ensemble learning methods and decision trees, Random Forest models demonstrate a good ability to capture intricate data interactions and patterns, thus providing a more precise depiction of the underlying trade dynamics when compared to linear gravity models.
- Graph Convolutional Network (GCN): A particular type of GNN intended for use with graph-structured data. By applying convolution-like processes to graphs, GCNs are able to capture the relationships and dependencies that exist between nodes. The following is the essential formula for a GCN layer's forward pass:

$$F^{(l+1)} = \sigma(D^{(-1/2)}AD^{(-1/2)}F^{(l)}W^{(l)}) \tag{2}$$

  where
  - $F^{(l)}$ is the node feature at layer l.
  - A is the adjacency matrix of the graph.
  - D is the degree matrix.
  - $W^{(l)}$ represents the weight parameters.
  - σ is the activation function.

  Each country is represented as a node within a graph, with initial feature vectors attached to these nodes serving as the foundation for developing node embeddings in the context of GCN, as detailed in Algorithm 1. These preliminary properties could include node attributes, such as GDP, trade value, distance, or other pertinent information. The primary concept in GCN is the propagation of information from neighboring nodes to update the embeddings based on Equation (2). The node embeddings $F^{(l+1)}$ are updated by aggregating and transforming information from neighboring nodes. Since the model modifies the weights $W^{(l)}$ during training to capture the relationships in the graph, these embeddings are learnable.

Algorithm 1 GCN Regression Model for Trade Value Prediction

Require: Train_Dataset
Require: country_code_map: Mapping of country codes to indices
Require: normalized_features: Array of features for each node
Require: normalized_target: Trade value for regression
Construct a DGL graph g
Assign node features to g
Define RegressionModel with GCN layers
Add self-loops to the graph g
Initialize the RegressionModel
Prepare for training: batch_size, optimizer, learning rate, epochs
for each epoch **do**
    for each batch in num_batches **do**
        Extract subgraph for batch nodes
        Compute logits using the model and batch's features
        Compute loss between logits and the batch's targets
        Perform backpropagation and optimizer step
        Accumulate epoch loss
    end for
    Record and print epoch loss
end for

- Graph Attention Network (GAT): A sophisticated variation of GNN that employs multiple attention methods to determine how important neighboring nodes are. The following formula is crucial for a GAT layer's forward pass:

$$f_i^{(l+1)} = \sigma\Big(\sum_{j \in N(i)} \alpha_{ij} W^{(l)} f_j^{(l)}\Big) \tag{3}$$

  where
  - $f_i^{(l+1)}$ is the updated node feature of node i at layer $l+1$.
  - α_{ij} represents the attention coefficient between nodes i and j.
  - $W^{(l)}$ is the weight matrix for the layer.
  - σ is the activation function.

  Bridging the gap between the theoretical underpinnings of GAT and their practical application, it becomes evident how Algorithm 2 harnesses these principles to transform international trade data analysis. The process begins with the development of node embeddings, which are crucial for capturing the intricacies of trade networks. GAT's innovative approach allows for the learning of node embeddings, utilizing attention techniques to prioritize neighboring nodes differently during aggregation. The attention mechanism establishes the relative relevance of each neighbor, and each node's embedding is calculated as a weighted sum of those embeddings based on Equation (3). The attention coefficients α_{ij}, introduced in GAT, indicate how much attention each neighbor node should receive throughout the aggregation and are calculated using a softmax function. During training, these embeddings, which are learnable parameters, are modified to reflect attention-weighted associations.

Algorithm 2 GAT Regression Model for Predicting Trade Values

Require: Train_Dataset
Require: country_code_map: mapping of country codes to indices
Require: normalized_features: Array of features for each node
Require: normalized_target: Trade value for regression
Create a DGL graph g
Assign normalized features to nodes of g.
Define the GAT regression model leveraging the attention mechanism.
procedure GATREGRESSIONMODEL(in_feats, $hidden_feats$, num_heads)
    Initialize GATConv layers with specified dimensions and heads.
    Define the forward pass implementing the attention mechanism.
end procedure
Add self-loops to g to include self-attention.
Initialize the GAT model specifying input, hidden dimensions, and heads.
Define optimizer and loss function.
for each epoch **do**
    for each batch **do**
        Extract subgraph for current batch.
        Compute logits using the model with attention.
        Compute loss between logits and target values.
        Backpropagate and update model weights.
    end for
    Record and print epoch loss
end for

4. Results

In this section, we assess the effectiveness of the Random Forest and two variations of the GNN model in analyzing international trade flows with the enriched UN Comtrade dataset. We outline the experimental configurations employed and subsequently delve into the outcomes observed from these experiments. Through these evaluations, we aim to gain insights into the performance and applicability of the GNN model variants.

4.1. Experimental Setups

The hypothesis posits that using GCN and GAT variations of GNN can improve international trade value prediction compared to the machine learning-based RF model, as described in Section 3. We used a particular hyperparameter tuning process to fine-tune the RF model to enhance its predictive capabilities. After comprehensive experimentation, we determined the optimal configuration, setting hyperparameters to $n_estimators = 100$, criterion = 'poisson', and $max_depth = 40$. This fine-tuned RF model balances complexity and accuracy, ensuring robust performance in predicting bilateral trade values. We leveraged the DGL library (Deep Graph Library, Version: 1.1.3, https://github.com/dmlc/dgl/ (accessed on 18 December 2023)) for implementing the GNN models and utilizing scikit-learn (scikit-learn, Version: 1.3.2, https://scikit-learn.org/stable (accessed on 23 December 2023)) for the evaluation metrics. The hyperparameters, as shown in Table 3, outline a learning rate of 0.0001, 100 epochs, and the use of the Adam optimizer. In the first experimental scenario, 80% of the entire edges were designated as training while the remaining 20% were assigned as tests for comparison with various methods. In the second scenario, we fixed the period from 2017 to 2021 for all the edges to be used as training edges, and the edges from 2022 were earmarked for testing.

The plotted loss function in Figure 4 provides valuable insights into the optimization process of our GAT model during training. As depicted, the loss steadily decreases over successive iterations, indicating that our model is effectively learning and improving its predictive performance. The smooth convergence of the loss function suggests that our training procedure is appropriately configured, leading to stable and consistent model updates. This visualization reaffirms the effectiveness of our training methodology and

provides confidence in the reliability of our model's predictions. Similar to the GAT model, the GCN also exhibits a rapid decrease in MSE loss during the initial epochs (as shown in Figure 5), indicating efficient early learning. However, the GCN's loss curve fluctuates more than the GAT model, especially in the first 10 epochs. Despite these initial instabilities, the loss stabilizes at a low value and remains relatively constant, although with minor fluctuations, throughout the remaining epochs. This behavior suggests that while GCN is effective, it may be more sensitive to initial weight settings and learning rates. This would lead to less smooth convergence compared to GAT. Overall, both models are effective, with the GAT model having a slight edge in stability and is potentially better for generalization.

Table 3. Hyperparameter values used during the experiment.

Hyperparameter	Value
Learning rate	0.0001
Batch size	10000
Number of epochs	100
Optimizer	Adam
Number of hidden layers	32

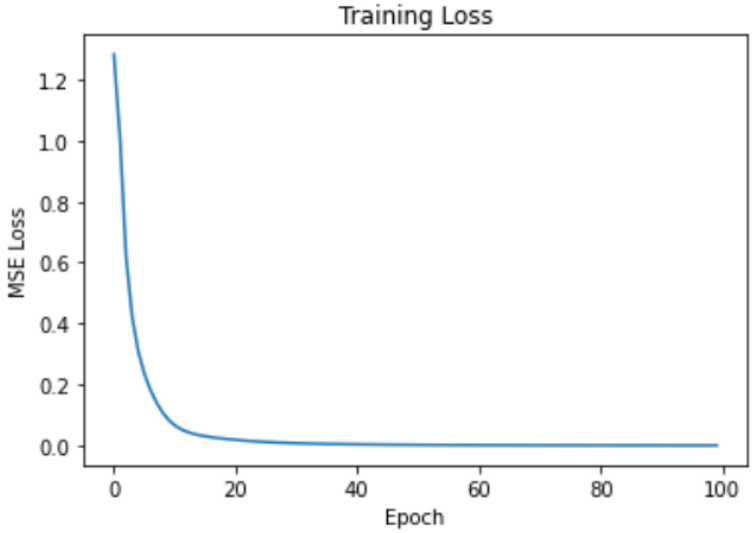


Figure 4. Loss function convergence during GAT training iterations.

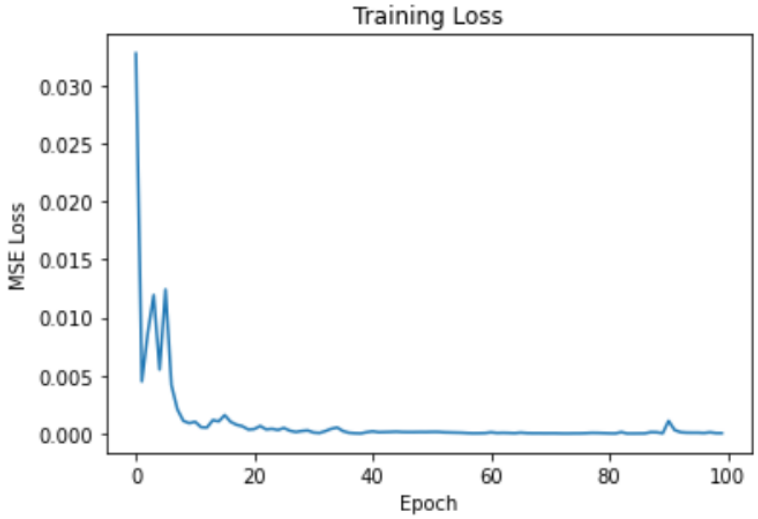


Figure 5. Loss function convergence during GCN training iterations.

4.2. Experimental Results

The experiment in this study marks a pivotal phase in our exploration of the predictive capabilities of three distinct models—RF, GCN, and GAT—within the realm of global trade analysis. To evaluate the performance of these models, this section examines various metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), R-squared values, Mean Absolute Percentage Error (MAPE), and Mean Percentage Error (MPE). Additionally, we utilize feature importances, which allows us to gain insights into the features influencing the model's predictions. Tables 4 and 5 present the performance of RF, GCN, and GAT models in predicting exports, using 2022 data for both seen and unseen scenarios.

In Table 4, the analysis uses two sets of features: one set includes CmdCode, Year, Distance, and GDP and the other decomposes GDP into GDP per capita and Population. This comparison aims to understand how well different machine learning models perform on the given dataset, and how additional economic and demographic indicators influence the models' predictive performance. The results show that GAT models, especially those with the expanded feature set, are more accurate in predicting export values, as evidenced by their lower MSE and MAE values and higher R-squared values. This indicates that GAT models are better at capturing the complex, non-linear relationships in the export data. However, the relatively high MAPE and MPE values for GAT models suggest a higher sensitivity to outliers, leading to larger errors in some cases. The RF model also performs well in terms of MAE, MAPE, and MPE, both with the standard data split and on unseen data. Still, GAT clearly outperforms RF and GCN in terms of R-squared value. The addition of GDP per capita and population significantly improves the GAT models' performance, increasing their R-square values and demonstrating a stronger correlation between predicted and actual export values. This underscores the value of incorporating broader economic and demographic contexts into predictive models for a more nuanced understanding of international trade dynamics.

Table 5 extends the analysis to evaluate the models' performance on unseen export data from 2022, using the same metrics and feature sets as in Table 4. The findings highlight the GCN models' superior performance on the unseen data, particularly when GDP per capita and population are included in the feature set. The MAPE and MPE values show that RF models tend to have lower error rates than GCN and GAT on unseen data. Managing error sensitivity, especially in the presence of outliers or extreme values, remains a challenge. The high MAPE and MPE values of GNN models indicate that while these models are generally accurate, they may occasionally produce significant errors.

Overall, the RF model demonstrates strong performance in terms of MAE, MAPE, and MPE. However, GAT outperforms RF in terms of MSE and R-squared value. The enhanced performance across all models with the inclusion of additional economic and demographic features on unseen data further supports the notion that a richer feature set can substantially improve a model's predictive accuracy and generalizability.

Table 4. Model performance comparison related to exports.

Model	Metrics	RF	GCN	GAT
CmdCode + Year + Distance + GDP	MSE	2994.09	565.97	**206.56**
	MAE	**1.71**	14.20	3.17
	MAPE	**1,585,289.78**	19,800,031.06	4,617,389.78
	MPE	**−1,585,270.71**	−8,423,804.78	−3,730,409.95
	R-square	0.60	0.67	**0.88**
CmdCode + Year + Distance + GDPcap + Population	MSE	2703.77	418.17	**153.09**
	MAE	**1.64**	5.78	3.22
	MAPE	**1,981,366.42**	9,330,344.70	4,007,344.39
	MPE	−1,981,347.43	−9,330,341.43	**−1,435,048.22**
	R-square	0.64	0.76	**0.91**

Table 5. Model performance comparison related to exports unseen 2022.

Model	Metrics	RF	GCN	GAT
CmdCode + Year + Distance + GDP	MSE	4591.19	1265.41	**223.20**
	MAE	2.13	26.32	**1.65**
	MAPE	2,265,935.22	76,888,108.94	**896,390.23**
	MPE	−2,265,915.08	−76,888,100.93	**−1,896,366.21**
	R-square	0.57	0.74	**0.95**
CmdCode + Year + Distance + GDPcap + Population	MSE	4234.66	387.26	**377.32**
	MAE	**2.12**	5.69	14.52
	MAPE	3,817,676.82	**17,083,815.17**	49,412,484.48
	MPE	−3,817,658.77	**−17,083,786.75**	−49,348,592.96
	R-square	0.60	**0.92**	**0.92**

5. Discussion

5.1. Performance of Graph Neural Networks

The research performed in this paper provides insight into how well three different models—Random Forest, Graph Convolutional Network, and Graph Attention Network—performed in predicting international trade patterns. Overall, the RF model exhibits robust performance according to certain criteria. Nevertheless, the GAT model surpasses RF in critical performance metrics such as MSE and R-squared value. The improvement in performance across all models, with the addition of extra economic and demographic features to unseen data, reinforces the idea that a more comprehensive set of features can significantly enhance a model's predictive accuracy and its ability to generalize.

From a machine learning perspective, the results highlight the efficacy of graph-based models, in capturing the complex relational dependencies inherent in trade data. The superior performance of GNNs can be attributed to their ability to leverage the graph structure of trade networks, allowing them to effectively learn from the relational information embedded in the data. The success of GNNs underscores the importance of incorporating graph-based approaches in predictive modeling tasks involving interconnected data entities, such as international trade networks.

Our research also identifies the pivotal role of input features in determining model performance. The feature importance analysis of RF, GCN, and GAT models reveals distinct priorities in predicting international trade flow. The RF model heavily relies on the commodity code (as shown in Figure 6), which highlights that the type of goods is the primary determinant, with significant but secondary importance placed on population sizes and geographic distance. This reflects RF's ability to capture global patterns across the dataset, indicating which features consistently contribute to making good splits across all trees.

In contrast, in Figure 7, the GCN model takes a more balanced approach by considering more features in the prediction process. The importance of features in GCNs often relates to how well they help aggregate information from neighboring nodes to enhance each node's representation. This indicates that GCN is effectively capturing the relational aspects of trade data, integrating economic and temporal factors to provide a holistic view. Notably, the living standards of the country of origin expressed through GDP per capita play a more important role than the country of origin in this model.

The GAT model provides the most balanced feature importance (as shown in Figure 8), integrating economic conditions, demographic factors, and temporal dynamics effectively. Features with higher attention weights in the GAT model are considered more relevant based on their neighboring nodes, making feature importance in GATs highly context-dependent and adaptable. This means that GAT is particularly effective at identifying and leveraging the contextual relationships within trade data. Consequently, it can more accurately predict instances where country is likely to trade with the neighbors of the existing trade partners, especially if those neighbors share similar economic indicators and demographic characteristics. Notably, the population of the country of origin is relatively

unimportant in informing the GAT model, suggesting that it can effectively handle both smaller and bigger countries.

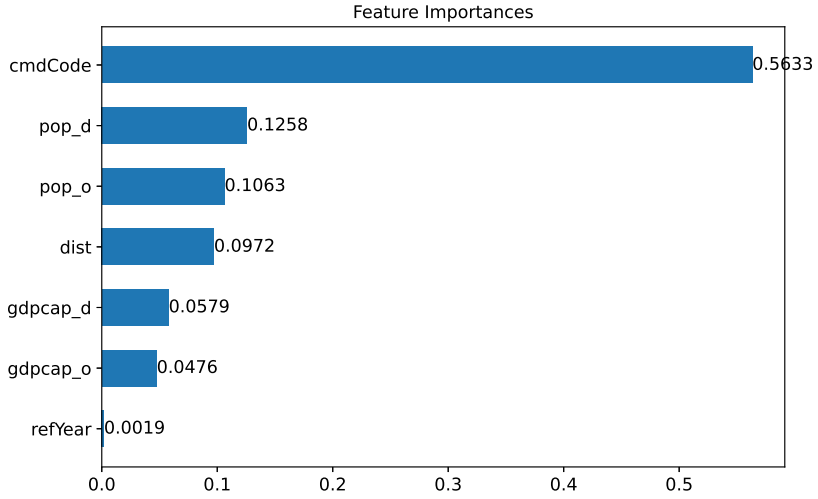


Figure 6. RF model feature importance.

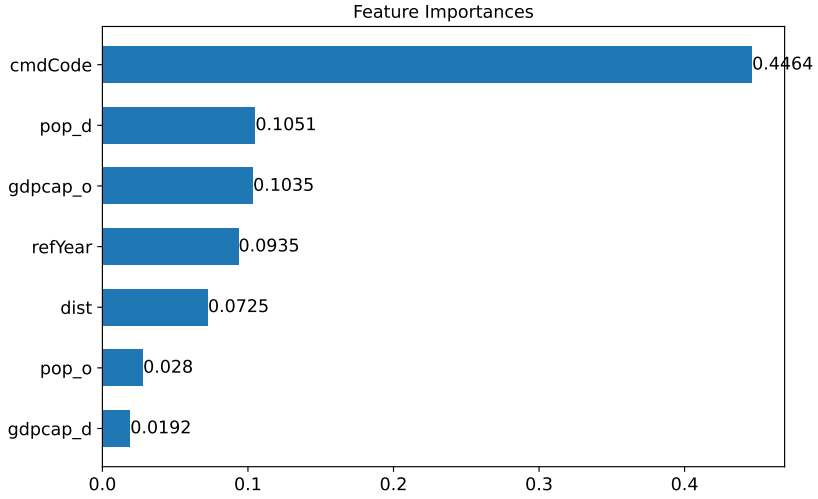


Figure 7. GCN model feature importance.

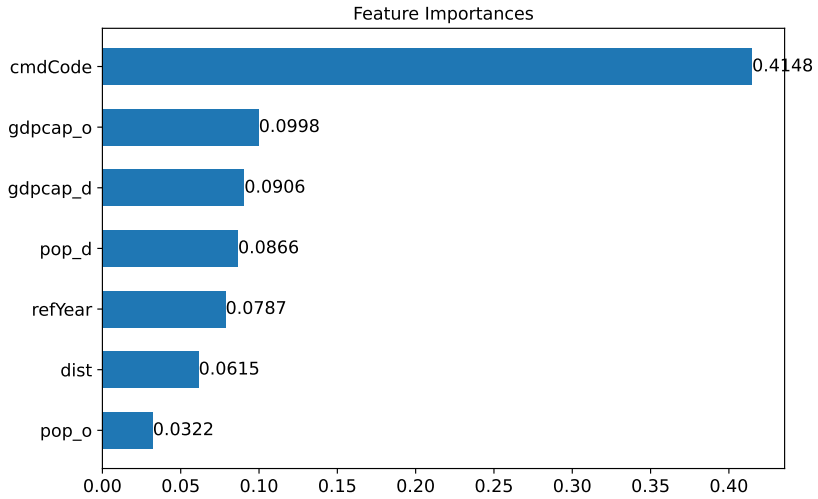


Figure 8. GAT model feature importance.

Overall, this analysis shows that RF, GCN, and GAT models capture different aspects of trade data; while RF offers a straightforward approach, GCN and GAT models leverage relational and contextual information, and GAT excels in capturing complex trade relationships by adjusting feature relevance based on neighboring nodes. These differences suggest that combining insights from RF and GNN models can lead to deeper data understanding and improved model performance through hybrid approaches and comprehensive feature engineering.

Furthermore, comparing seen and unseen data from 2022 underscores the importance of model robustness and generalization. Despite differences in performance across models and feature sets, consistent trends emerge, highlighting the reliability of the models in predicting trade dynamics across different periods.

The importance of features such as GDP, GDP per capita, and population reaffirms their role as fundamental determinants of trade patterns. Countries with higher economic output and larger populations tend to engage in more extensive trade networks, highlighting the importance of economic size and market potential in driving international trade. Additionally, the influence of distance on trade flows underscores the continued relevance of geographic proximity in shaping trade relations, despite advances in technology and globalization.

5.2. Data Drift Analysis

In the dynamic landscape of international trade analysis, machine learning models' sensitivity to data drift is a critical consideration [35]. Changes to consumer sentiment, modifications in trade policies leading to restructuring of value chains and geopolitical events are just a few examples of the factors that can cause changes in the statistical features of the input data. As a result, the Kolmogorov–Smirnov (KS) statistic [36] becomes useful in this situation. It quantifies the extent to which the current set of model outputs deviates from a reference dataset, typically a period of stability in the model's historical performance.

Its potential application involves a continuous monitoring system that regularly calculates the KS statistic, offering a quantitative measure of the extent of data drift. When the KS statistic goes above a certain level, indicating changes in data patterns, it starts an important response mechanism: the models are retrained. Retraining becomes imperative to ensure the models promptly adapt to the evolving dynamics, assimilating the most recent data and refining their predictive capabilities.

The KS statistic, acting as a dynamic sentinel, not only alerts to data drift but also guides the strategic refinement of the models. By embracing a data-driven retraining strategy, the models become more adept at capturing nuanced shifts in international trade patterns, thereby enhancing their accuracy and relevance over time. This all-around method emphasizes keeping the models running smoothly even when data changes.

In our case, we analyze the drift of our dataset used to train our RF, GCN, and GAT models. As illustrated in Figure 9, a comprehensive plot measures the KS statistic, providing a graphical depiction of the data drift experienced by the family of products over the years. This graphical representation not only enhances interpretability but also identifies distinct periods exhibiting notable drift. By scrutinizing the plotted KS values over time, we can pinpoint specific years where deviations in the data distribution occur, such as the period between 2019 and 2022, signifying potential shifts in the dynamics of international trade, specifically for some categories of products.

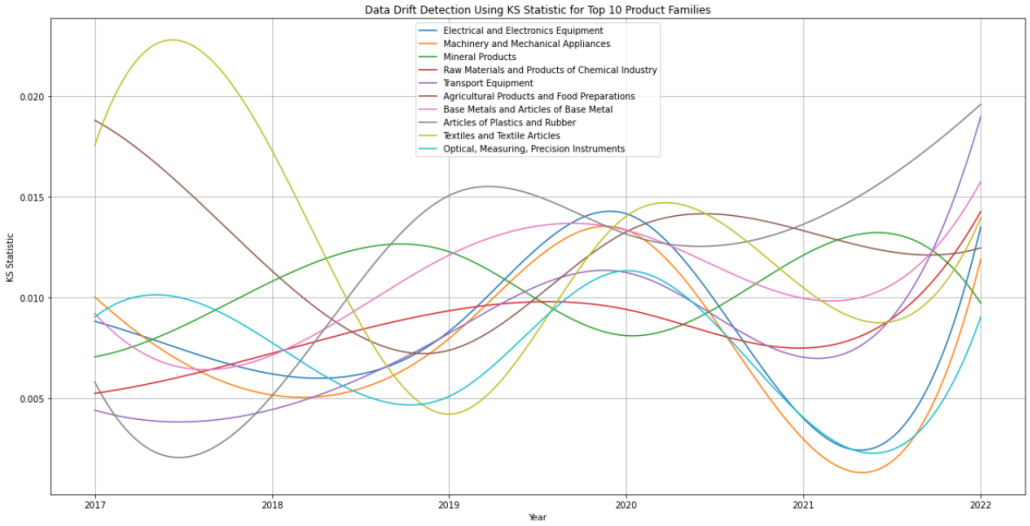


Figure 9. Data drift for products' family.

5.3. Relevance for Business and Policy Stakeholders

The fields of economic geography, international development, and innovation studies are currently benefiting from advancements using international trade data and rapidly evolving analytical methods [3,37]. These methods are instrumental in deciphering the complex relationships between spatial factors and economic development trajectories; they facilitate the identification of potential product categories for economic diversification, thereby equipping policymakers and development organizations with the tools to forecast and guide the direction of economic growth and to identify strategic investment areas that can promote sustainable and equitable development. These methods offer a solid framework to comprehend the spread of innovative ideas and technologies across geographies, shedding light on the subtleties of industry evolution, the emergence of new industries, and the transformation of regional economies [38,39].

By analyzing the potential export value of products, policymakers and entrepreneurs can identify opportunities that are poised to generate substantial export revenues. A detailed evaluation of potential market demand and accessibility is particularly vital for smaller economies that may not be competitive based on volume alone, making it essential to target high-value niches within the global market [40,41]. In this scenario, discerning which products are likely to succeed in foreign markets and the value that they could achieve overseas enables these economies to focus on products with promising prospects. Accurate predictions of new product opportunities and bilateral export potential are crucial for entrepreneurs and policymakers in small to medium-sized economies striving to foster sustainable economic growth [7,8].

To showcase the broad applicability of the proposed graph models, we selected Estonia, Brazil, India, and Malaysia as case study countries. These nations exhibit diverse sizes and geographical locations across different continents. Using the best-performing GAT model, we analyzed their exporting opportunities for different products in 2022.

For example, the GAT model, in Figure 10, indicates that Estonia's predicted exports of medicaments are concentrated in specific regions, with the highest volumes expected in Finland, Belgium, and India. There is also a notable presence elsewhere in Europe. The varying shades of green highlight the differing levels of market demand across the globe. This suggests Estonia's medicament export strategy is likely focused on maintaining and expanding in high-demand regions while exploring potential growth in moderately predicted areas. The currently minimal presence in other parts of the world might point to opportunities for market expansion or reflect current trade limitations.

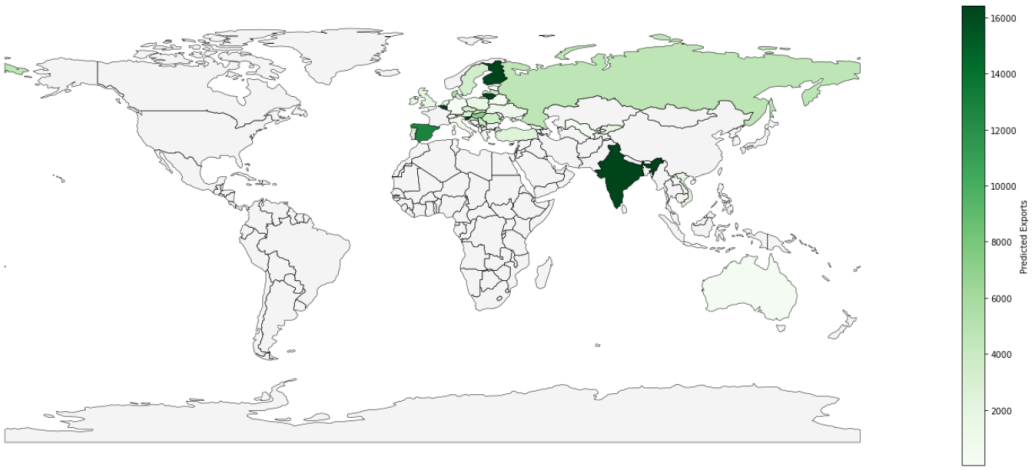


Figure 10. Estonia's predicted export of medicaments (HS 300490).

For Brazil, the predicted export markets in medical instruments in Figure 11 conform to expected market dynamics, indicating strong demand in Russia, South Africa, and Argentina, along with a substantial presence in North America, various Asian regions, and several African nations. However, the absence of predicted exports to key markets such as parts of Europe and the Middle East highlights some of the avenues for further analysis.

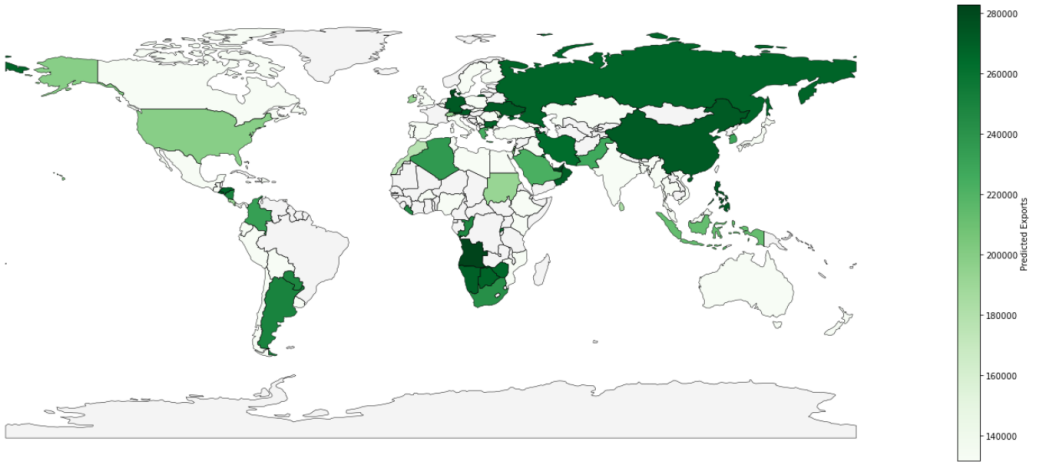


Figure 11. Brazil's predicted export of medical instruments (HS 901890).

India's predictions for laptop computers generally align with industry life cycles, indicating high demand in North America, parts of Africa, and some Asian countries (as shown in Figure 12). However, the absence of predicted exports to certain key markets raises questions about potential gaps or assumptions in the model.

The analysis employing the GAT model for Malaysia's medicaments indicates robust predicted exports across numerous global markets (as shown in Figure 13). As of 2022, 14 countries accounted for 81% of Malaysia's exports in this product group [42], suggesting for substantial opportunities for diversification of export markets.

Still, these small case studies should be approached with caution. It is not always practical to base strategic development solely on market demand predictions from the aforementioned models. To improve the effectiveness of economic and business development strategies, future research should include a wider array of variables such as trade barriers, market concentration, expected market demand growth rates, and the trade risks associated with destination markets [43]. These elements are pivotal in influencing trade flows and, by extension, the success of export initiatives. A practical approach to overcoming these challenges involves a more comprehensive synthesis of international trade research with global marketing strategies [3,42,44–47]. Consequently, future research should explore

additional factors and integrate dynamic elements like trade agreements, market demand growth rates, exchange rates, political stability, geopolitical distance, and technological advancements to enhance the predictive accuracy and practical relevance of these models.

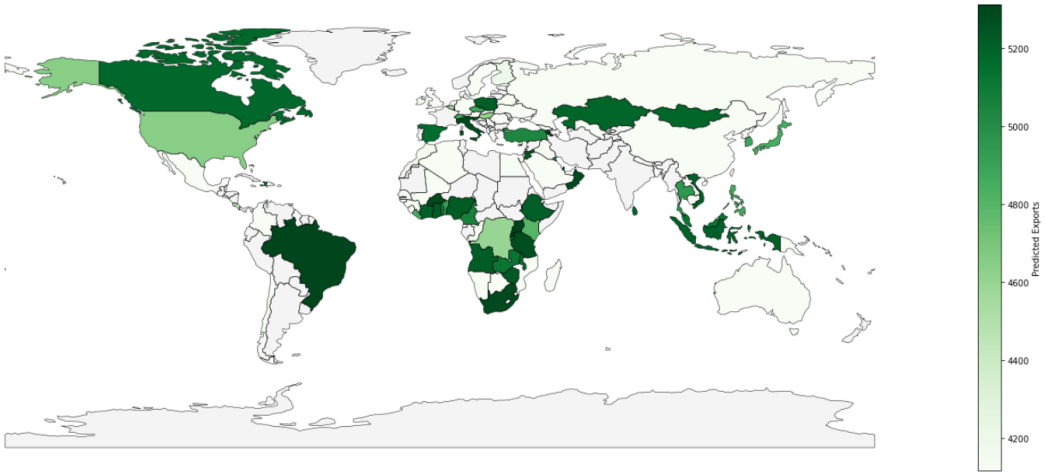


Figure 12. India's predicted export in laptop computers (HS 847130).

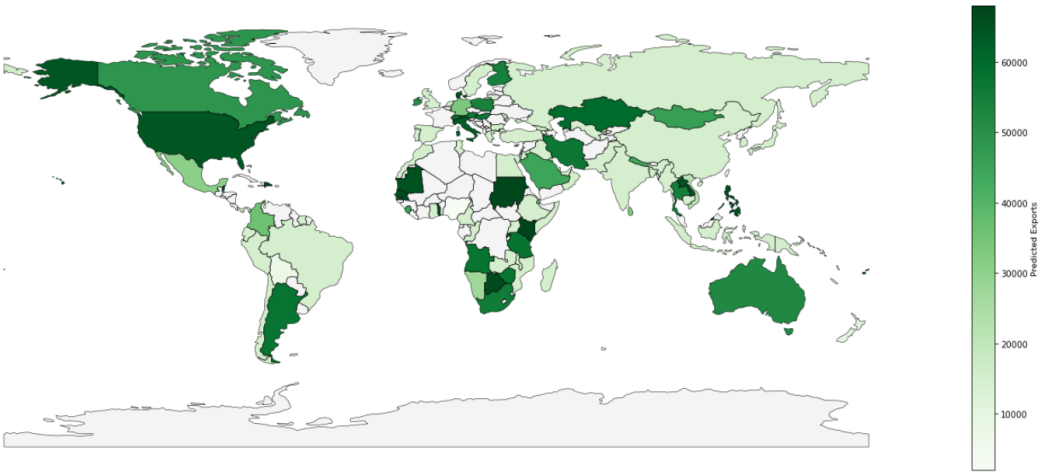


Figure 13. Malaysia's predicted export in medicaments (HS 300490).

6. Conclusions and Outlooks

The use of machine learning techniques in estimating international trade marks a considerable change, providing new opportunities compared to the conventional regression models that have predominantly governed the field. In the past, computational limitations and strict methodological frameworks had limited international trade prediction models, leading them to rely heavily on aggregated data or concentrate narrowly on a limited range of product categories. This study pioneers applying machine learning to handle the complexities involved in analyzing thousands of product groups within a single model. This novel approach and the increase in the available computing power not only overcome previous computational and methodological limitations but also deliver a level of granularity and accuracy in trade predictions that were previously unattainable.

This article's primary contribution is the demonstration that machine learning methods can handle thousands of product groups in one integrated model, predicting bilateral trade flows at the detailed six-digit HS product code level. We leverage advanced AI models, specifically GNN, to predict trade values among nations. Based on a comparison study with the RF model, we have demonstrated that GNN models outperform RF models in terms of predictive power. We have highlighted the inherent sensitivity of these models

to data drift, a phenomenon that can significantly impact their accuracy over time. By employing the Kolmogorov–Smirnov (KS) statistic and visually representing data drift for the family of products over the years, we have not only quantified but also visualized the dynamic changes in data distribution.

Our findings not only emphasize the superiority of GNN models in capturing the complexities of international trade but also highlight the significance of understanding and addressing data drift for sustained model accuracy. Our results underscore the importance of proactive strategies, including continuous monitoring and timely model retraining, to mitigate the adverse effects of data drift. The KS statistic is a strong indicator that helps us understand the subtleties of temporal patterns in data drift, guides interventions, and makes sure that AI models can adapt to how international trade is changing.

As future work, we could focus on further improving the resilience and responsiveness of forecasting models. Further research could investigate the integration of additional metrics to comprehensively assess data drift, taking into account the multifaceted nature of international trade dynamics. Moreover, investigating set methods or hybrid models that combine GNNs with traditional machine learning approaches could provide a more holistic and resilient forecasting framework. We could also make forecasting models that are even more accurate and responsive by looking into ways to adapt to real-time data drift and how events in other countries affect trade dynamics. This study establishes the foundation for further research that seeks to improve the precision, flexibility, and dependability of AI models in the constantly changing context of global trade.

Author Contributions: Conceptualization, B.S., C.O., T.K., M.T. and D.R.-Y.; formal analysis, B.S. and C.O.; investigation, M.T. and D.R.-Y.; methodology, T.K. and D.R.-Y.; resources, C.O.; software, B.S.; supervision, T.K. and M.T.; validation, T.K. and M.T.; visualization, B.S.; writing—original draft, B.S., C.O., M.T. and D.R.-Y.; writing—review and editing, T.K. and M.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received support from grants awarded to TalTech by the Estonian Research Council for the project titled "Economic Complexity, Machine Learning, and Economic Policy" (PRG1573), as well as from the TalTech Industrial project (952410) and the CatChain project (778398) under the European Horizon 2020 programme.

Data Availability Statement: The data that support the findings of this study are not publicly available due to licensing constraints imposed by the United Nations Statistical Division.

Acknowledgments: The authors express their gratitude to the editors and reviewers for their valuable comments and constructive suggestions regarding the revision of this paper.

Conflicts of Interest: All authors hereby declare that there are no conflicts of interest regarding the data and manuscript.

References

1. Isard, W. Location Theory and Trade Theory: Short-Run Analysis. *Q. J. Econ.* **1954**, *68*, 305–320. [CrossRef]
2. Tinbergen, J. *Shaping the World Economy: Suggestions for an International Economic Policy*; Twentieth Century Fund: New York, NY, USA, 1962. [CrossRef]
3. Tiits, M.; Kalvet, T.; Ounoughi, C.; Ben Yahia, S. Relatedness and Product Complexity Meet Gravity Models of International Trade. *J. Open Innov. Technol. Mark. Complex.* **2024**, *10*, 100288. [CrossRef]
4. Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; Yu, P.S. A Comprehensive Survey on Graph Neural Networks. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *32*, 4–24. [CrossRef] [PubMed]
5. Wu, L.; Cui, P.; Pei, J.; Zhao, L.; Guo, X. Graph Neural Networks: Foundation, Frontiers, and Applications. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 14–18 August 2022; Association for Computing Machinery: New York, NY, USA, 2022; pp. 4840–4841. [CrossRef]
6. Skarding, J.; Gabrys, B.; Musial, K. Foundations and Modeling of Dynamic Networks Using Dynamic Graph Neural Networks: A Survey. *IEEE Access* **2021**, *9*, 79143–79168. [CrossRef]
7. Tiits, M.; Kalvet, T. Intelligent piggybacking: A foresight policy tool for small catching-up economies. *Int. J. Foresight Innov. Policy* **2013**, *9*, 253–268. [CrossRef]
8. Ploom, I.; Kalvet, T.; Tiits, M. Defence industries in small European states: Key contemporary challenges and opportunities. *J. Int. Stud.* **2022**, *15*, 112–130. [CrossRef]

9. Capoani, L. Review of the gravity model: Origins and critical analysis of its theoretical development. *SN Bus. Econ.* **2023**, *3*, 95. [CrossRef]
10. Sharma, P.; Rohatgi, S.; Jasuja, D. Scientific Mapping of Gravity Model of International Trade Literature: A Bibliometric Analysis. *J. Scientometr. Res.* **2023**, *11*, 447–457. [CrossRef]
11. Hillberry, R.; Hummels, D. Intranational Home Bias: Some Explanations. *Rev. Econ. Stat.* **2003**, *85*, 1089–1092. [CrossRef]
12. Yotov, Y.V.; Piermartini, R.; Monteiro, J.-A.; Larch, M. *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*; WTO iLibrary: Genève, Switzerland, 2016.
13. Anderson, J.E.; Marcouiller, D. Insecurity and the Pattern of Trade: An Empirical Investigation. *Rev. Econ. Stat.* **2002**, *84*, 342–352. http://www.jstor.org/stable/3211782 [CrossRef]
14. Rincon-Yanez, D.; Mouakher, A.; Senatore, S. Enhancing downstream tasks in Knowledge Graphs Embeddings: A Complement Graph-based Approach Applied to Bilateral Trade. *Procedia Comput. Sci.* **2023**, **225**, 3692–3700. [CrossRef]
15. Baldwin, R.; Taglioni, D. *Gravity for Dummies and Dummies for Gravity Equations*; Working Paper; National Bureau of Economic Research: Cambridge, UK, 2006. [CrossRef]
16. Jun, B.; Alshamsi, A.; Gao, J.; Hidalgo, C.A. Bilateral relatedness: Knowledge diffusion and the evolution of bilateral trade. *J. Evol. Econ.* **2020**, *30*, 247–277. [CrossRef]
17. Athey, S.; Imbens, G. Machine Learning Methods That Economists Should Know About. *Annu. Rev. Econ.* **2019**, *11*, 685–725. [CrossRef]
18. James, G.; Witten, D.; Hastie, T.; Tibshirani, R.; Taylor, J. *An Introduction to Statistical Learning: With Applications in Python*; Springer: Berlin/Heidelberg, Germany, 2023. [CrossRef]
19. Ho, T. Random decision forests. In Proceedings of the 3rd International Conference on Document Analysis and Recognition, Montreal, QC, Canada, 14–16 August 1995; Volume 1, pp. 278–282.
20. Scarselli, F.; Gori, M.; Tsoi, A.; Hagenbuchner, M.; Monfardini, G. The graph neural network model. *IEEE Trans. Neural Netw.* **2008**, *20*, 61–80. [CrossRef]
21. Li, J.; Rong, Y.; Cheng, H.; Meng, H.; Huang, W.; Huang, J. Semi-supervised graph classification: A hierarchical graph perspective. In Proceedings of the World Wide Web Conference, San Francisco, CA, USA, 13–17 May 2019; pp. 972–982. [CrossRef]
22. Kipf, T.N.; Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv* **2016**, arXiv:1609.02907. https://doi.org/10.48550/arXiv.1609.02907.
23. Li, Y.; Yu, R.; Shahabi, C.; Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv* **2017**, arXiv:1707.01926. https://doi.org/10.48550/arXiv.1707.01926.
24. Verstyuk, S.; Douglas, M. Machine Learning the Gravity Equation for International Trade. 2022. Available online: https://ssrn.com/abstract=4053795 (accessed on 8 March 2024).
25. Minakawa, N.; Izumi, K.; Sakaji, H. Bilateral Trade Flow Prediction by Gravity-informed Graph Auto-encoder. In Proceedings of the 2022 IEEE International Conference On Big Data (Big Data), Osaka, Japan, 17–20 December 2022; pp. 2327–2332. [CrossRef]
26. Monken, A.; Haberkorn, F.; Gopinath, M.; Freeman, L.; Batarseh, F. Graph neural networks for modeling causality in international trade. *Int. Flairs Conf. Proc.* **2021**, *34*. [CrossRef]
27. Atwood, J.; Towsley, D. Diffusion-convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **2016**, *29*. [CrossRef]
28. Cao, S.; Lu, W.; Xu, Q. Deep neural networks for learning graph representations. In Proceedings of the AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; Volume 30. [CrossRef]
29. Qiu, J.; Dong, Y.; Ma, H.; Li, J.; Wang, K.; Tang, J. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, Los Angeles, CA, USA, 5–9 February 2018; pp. 459–467. [CrossRef]
30. Rincon-Yanez, D.; Ounoughi, C.; Sellami, B.; Kalvet, T.; Tiits, M.; Senatore, S.; Yahia, S. Accurate prediction of international trade flows: Leveraging knowledge graphs and their embeddings. *J. King Saud Univ.-Comput. Inf. Sci.* **2023**, *35* , 101789. [CrossRef]
31. Ahmed, F.; Cui, Y.; Fu, Y.; Chen, W. A graph neural network approach for product relationship prediction. In Proceedings of the International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, Online, 17–19 August 2021; Volume 85383, p. V03AT03A036. [CrossRef]
32. Panford-Quainoo, K.; Bose, A.; Defferrard, M. Bilateral Trade Modelling with Graph Neural Networks. ICLR Workshop on Practical ML for Developing Countries. 2020. Available online: https://www.researchgate.net/profile/Kobby-Panford-Quainoo/publication/339200492_Bilateral_Trade_Modeling_with_Graph_Neural_Networks/links/5f781d37299bf1b53e099940/Bilateral-Trade-Modeling-with-Graph-Neural-Networks.pdf (accessed on 22 May 2024).
33. Di Paolo, G.; Rincon-Yanez, D.; Senatore, S. A Quick Prototype for Assessing OpenIE Knowledge Graph-Based Question-Answering Systems. *Information* **2023**, *14*, 186. [CrossRef]
34. Rincon-Yanez, D.; Senatore, S. FAIR Knowledge Graph construction from text, an approach applied to fictional novels. In Proceedings of the 1st International Workshop on Knowledge Graph Generation from Text and the 1st International Workshop on Modular Knowledge Co-Located with 19th Extended Semantic Conference (ESWC 2022), Crete, Greece, 29 May–2 June 2022; pp. 94–108.
35. Poenaru-Olaru, L.; Cruz, L.; Rellermeyer, J.; Van Deursen, A. Maintaining and monitoring AIOps models against concept drift. In Proceedings of the 2023 IEEE/ACM 2nd International Conference on AI Engineering–Software Engineering for AI (CAIN), Melbourne, Australia, 15–16 May 2023; pp. 98–99.

36. Massey, F., Jr. The Kolmogorov-Smirnov test for goodness of fit. *J. Am. Stat. Assoc.* **1951**, *46*, 68–78. [CrossRef]
37. Hidalgo, C.A. Economic complexity theory and applications. *Nat. Rev. Phys.* **2021**, *3*, 92–113. [CrossRef]
38. Balland, P.-A.; Jara-Figueroa, C.; Petralia, S.; Steijn, M.; Rigby, D.L.; Hidalgo, C.A. Complex economic activities concentrate in large cities. *Nat. Hum. Behav.* **2022**, *6*, 435–446. [CrossRef] [PubMed]
39. Tiits, M.; Karo, E.; Kalvet, T. Small countries facing the technological revolution: Fostering synergies between economic complexity and foresight research. *Compet. Rev.* 2024, *in press*. [CrossRef]
40. Kattel, R.; Randma-Liiv, T.; Kalvet, T. Small States, Innovation and Administrative Capacity. In *Innovation in the Public Sector: Linking Capacity and Leadership*; Bekkers, V., Edelenbos, J., Steijn, B., Eds.; Palgrave Macmillan: London, UK, 2011; pp. 61–81. [CrossRef]
41. Tiits, M; Kalvet, T. *Nordic Small Countries in the Global High-Tech Value Chains: The Case of Telecommunications Systems Production in Estonia*; Working Papers in Technology Governance and Economic Dynamics; TUT Ragnar Nurkse Department of Innovation and Governance: Tallinn, Estonia, 2012. Available online: http://technologygovernance.eu/files/main/2012022211372121.pdf (accessed on 22 May 2024).
42. Tiits, M.; Kalvet, T.; Mehide, I. *Goodtrade.ai Export Strategy Analytics Platform*; Policy Lab: Brooklyn, NY, USA, 2023. Available online: https://www.goodtrade.ai/ (accessed on 22 May 2024).
43. Kalvet, T.; Tiits, M.; Ounoughi, C.; Ben Sassi, I.; Ben Yahia, S. At the Crossroads of Product Complexity, Market Demand, and Machine Learning. *Manag. Mark. J.* 2024, *accepted for publication*.
44. Cuyvers, L.; Viviers, W. (Eds.) *Export Promotion: A Decision Support Model Approach*; Sun Press: Audubon, NJ, USA, 2012.
45. Cameron, M.; Cuyvers, L.J.; Fu, D.; Viviers, W. Identifying export opportunities for China in the "Belt and Road Initiative" group of countries: A decision support model approach. *J. Int. Trade Law Policy* **2021**, *20*, 101–126. [CrossRef]
46. Aucamp, M.; Steenkamp, E.A.; Bezuidenhout, C. Comparing International Market Selection Methods Using Export Potential Values for South Africa. *Int. Trade J.* **2023**, 1–23. [CrossRef]
47. Kalvet, T.; Tiits, M. Identification of Export-led Catching-up Opportunities in Turbulent Times. 2024; Unpublished manuscript.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.