

# Graph Neural Networks for Modeling Causality in International Trade

Anderson Monken<sup>\*1</sup>, Flora Haberkorn<sup>\*1</sup>, Munisamy Gopinath<sup>2</sup>, Laura Freeman<sup>3</sup>, Feras A. Batarseh<sup>3</sup>

<sup>1</sup>International Finance Division of the Federal Reserve Board, Washington D.C., USA

(anderson.e.monken@frb.gov, flora.m.haberkorn@frb.gov)

<sup>2</sup>University of Georgia, Athens, GA, USA (m.gopinath@uga.edu)

<sup>3</sup>Virginia Polytechnic Institute and State University (Virginia Tech), Arlington, VA, USA

(batarseh@vt.edu, laura.freeman@vt.edu)

## Abstract

Neural network algorithms have proven successful for accurate classifications in many domains such as image recognition and semantic parsing. However, they have long suffered from the lack of ability to measure causality, predict outliers effectively, or provide explainability relevant to the application domain. In this paper we introduce a method that measures causal scenarios during outlier events using neural networks: Artificial Intelligence Network Explanation of Trade (AINET). AINET tailors AI techniques specifically for bilateral trade modeling. Datasets with network-like structures (such as global trade, social networks, or city traffic) can benefit from Graph Neural Networks (GNNs) modeling and structural power. These network-based models (i.e. GNNs) empower policy makers with an understanding of the fast-paced shifts in trade flows around the world due to outlier events such as increased tariffs, natural disasters, embargoes, pandemics, or trade wars. Our work is at the intersection of GNNs' optimization, causality, and their proper application to trade. AINET results are presented with an overall test mean absolute percentage error (MAPE) of 28%, demonstrating the efficacy and potential of harnessing this method.

## Introduction and Motivation

A key objective of quantitative economic analyses is to uncover relationships – e.g. supply, prices, or trade – for use in making predictions or forecasts of future outcomes. Current systems that generate forecasts for decision making tend to use ad hoc, expert-driven, linear, and non-linear models or Bayesian approaches. AI methods, however, provide solutions to causality and relationships while mitigating imbalance, bias, and outliers with ease.

Many sub-fields of economics are moving towards the use of AI in behavioral modeling, e.g. (Ghoddusi, Creamer, and Rafizadeh 2019) for energy economics. However, there is limited work in the literature that shows tailored and optimized neural network algorithms are effective for economic

applications. Neural networks offer the ability to find connections between variables or features that are difficult to detect using econometrics (traditional methods). Even compared to other machine learning methods, neural networks can provide a better fit in the long term as found in (Gopinath et al. 2021) for agricultural trade. The 2020 U.S. National Artificial Intelligence Initiative Bill states: “Artificial intelligence is a tool that has the potential to change and possibly transform every sector of the U.S. economy” (S. 1558 - AIIA 2020). Given its obvious influence on the economy, the next sector for AI to advance is international trade, and its causal effects on production, consumption, and prices.

Understanding trade flows has historically been a challenge, particularly in the context of articulating projections. Not only are trade flows influenced heavily by trade partner policies, but also by exogenous factors and unanticipated shocks. The global soybean market is a case in point, since it is the largest agricultural commodity traded and commonly targeted during international disputes. The biggest exporters of soybeans are Brazil and the U.S., while the biggest importers are China and the European Union (E.U.), with China importing ten times as much as the E.U.

The complexity of international trade naturally makes it a network, thus, utilizing a network model lets us map relationships and conduct experimental counterfactual scenarios. In our work, modeling soybeans using GNNs illustrates the need for tailoring and applying advanced AI algorithms to study complex global commodity flows. Moreover, highlighting this crucial commodity provides a glimpse into the global impact consequences of trade wars and pandemics on agricultural goods.

## Related Works in Neural Networks

Understanding, modeling, and predicting international trade is a central domain in economics. Economists have implemented the *Gravity* equation of trade to describe trade flows (Anderson 1979). More recently, researchers across multi-disciplinary fields have identified international trade as well suited for network analysis by considering the breadth/value of trade between countries as weighted edges, and countries as nodes (Bhattacharya, Mukherjee, and Manna 2007). By harnessing machine learning methods, researchers have been able to better predict agricultural exports trade patterns using Gradient Tree Boosting, ARIMA, and XGBoosting

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

\*The views expressed in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System.

(Batarseh et al. 2019). Use of *shallow* neural networks also improves international trade forecasting (Wohl and Kennedy 2018). Combining neural networks with network analysis theory of international trade, (Panford-Quainoo, Bose, and Defferrard 2020) demonstrates the efficacy of country GDP classification using a GNN for a single time period. Our contribution to the literature is tuning GNNs in novel ways (with time-varying network structures) to model trade flows for predicting future patterns under unforeseen circumstances, understanding causality of such events by remodeling the outcomes in novel ways, and influencing public policy for the greater good.

While grappling with new methods and uncertainties, the availability of big data and advances in software systems pose new challenges to conventional approaches to forecasting. These include dealing with the sheer volume of data, the lengthy list of variables/features available to explain such relationships (and associated collinearity issues), and the need to move beyond failed static and linear models. Our GNNs are offered as an alternative to address many of these challenges in economic modeling (Mullainathan and Spiess 2017; Batarseh and Yang 2017).

Convolutional Neural Networks (CNNs) are popular for many domains including textual analysis and computer vision (Kim 2014). CNNs apply a filter with shared weights over input data and potentially pool the results to reduce dimensions, creating a local receptive field where information is aggregated in the neighborhood of the surrounding words or pixels (Yamaguchi et al. 1990; Lecun et al. 1998). Equation 1 shows maximum pooling, which reduces the dimension of the output tensor by grabbing the max value using a kernel size  $k$  and an optional stride to skip pixels  $stride$ . The resulting output of a series of the convolution and pooling layers can then be an image classification or a binary economic output:

$$out(N_i, C_j, h, w) = \max_{m=0, \dots, kH-1} \max_{n=0, \dots, kW-1} input(N_i, C_j, stride[0] \times h + m, stride[1] \times w + n) \quad (1)$$

Conversely, Graph Convolutional Networks (GCNs) generalize the convolution process by aggregating information among nodes rather than scanning over a grid of pixels, creating a neighborhood effect so that nodes with fewer degrees of separation and stronger weighted edges tend to have more common features. The Chebyshev filter size,  $K$ , determines the sphere of influence between the node of interest and the neighborhood of nodes when convolving and pooling information for models (Defferrard, Bresson, and Vandergheynst 2017). A  $K$  of 1 would not include neighborhood effects, while a  $K$  of 2 would take into account nodes with a direct edge with the node of interest.

Recurrent Neural Networks (RNNs) take into account the sequential connection of data records, which is heavily relevant to trade flows, allowing for an observation at time  $t-1$  to influence predictions at time  $t$  (Elman 1990). Equation 2 is a simple RNN framework with  $W_{ih}$  the weight matrix for input data  $x_t$ , two bias terms  $b_{ih}$  and  $b_{hh}$ , and the weight matrix  $W_{hh}$  for the previous hidden state matrix output  $h_{t-1}$ . The sum then undergoes a non-linear transformation, such

as  $\tanh$ , to create the hidden matrix output  $h_t$ .

$$h_t = \tanh(W_{ih}x_t + b_{ih} + W_{hh}h_{t-1} + b_{hh}) \quad (2)$$

In the case of causality and outliers' evaluation for trade flows, we aim for the model to learn from all data regardless of timestamps or geographies; while neural networks have issues associated with learning from subtle events, Long-Short Term Memory (LSTM) allows for the resolution of that challenge (discussed further in the assurance section of this paper). LSTMs include both a cell state  $c_{t-1}$  and hidden state  $h_{t-1}$  to create logic gates which selectively pass information over time (Hochreiter and Schmidhuber 1997). The LSTM framework was extended to a fully-connected version to enable long-term structural connections, and convolutions were incorporated into LSTMs to allow for spatial-temporal analysis (Graves 2014; Shi et al. 2015). Equation 3 shows a convolutional LSTM cell.

$$\begin{aligned} i &= \sigma(W_{xi} * x_t + W_{hi} * h_{t-1} + w_{ci} \odot c_{t-1} + b_i) \\ f &= \sigma(W_{xf} * x_t + W_{hf} * h_{t-1} + w_{cf} \odot c_{t-1} + b_f) \\ c_t &= f_t \odot c_{t-1} + i_t \odot \tanh(W_{xc} * x_t + W_{hc} * h_{t-1} + b_c) \\ o &= \sigma(W_{xo} * x_t + W_{ho} * h_{t-1} + w_{co} \odot c_{t-1} + b_o) \\ h_t &= o_t \odot \tanh(c_t) \end{aligned} \quad (3)$$

This model has an input gate  $i_t$ , forget gate  $f_t$ , cell gate  $c_t$ , and output gate  $o$  to yield the final  $h_t$  output matrix (Shi et al. 2015). Each gate includes convolutions of the input data.

## Methods: AINET

AINET is a method for analyzing the time-varying structure of the network of bilateral country trade using GNNs. AINET is based on a stateless recurrent network model with underlying Chebyshev filtered convolutions. See Equation 5.

GCNs rely on capturing the dynamics of a static graph of data, though changing the adjacency matrix of the network reduces the effectiveness of the model. Spatial-temporal methods of prediction using recurrent GNNs combine the graph-modified convolution and the flexibility of a model that can adapt to changes over time. Combining elements of RNNs and GNNs *specifically for our trade deployment* increases the effective model capacity in graph convolutional LSTM models (GC-LSTM), as shown in Equation 4. The  $x_t$  and  $h_{t-1}$  convolutions are replaced by graph convolutions, which harness Chebyshev filtering. A notion that could be critical to the success of neural networks in other domains as well.

$$\begin{aligned} i &= \sigma(W_{xi} *_{\mathcal{G}} x_t + W_{hi} *_{\mathcal{G}} h_{t-1} + w_{ci} \odot c_{t-1} + b_i) \\ f &= \sigma(W_{xf} *_{\mathcal{G}} x_t + W_{hf} *_{\mathcal{G}} h_{t-1} + w_{cf} \odot c_{t-1} + b_f) \\ c_t &= f_t \odot c_{t-1} + i_t \odot \tanh(W_{xc} *_{\mathcal{G}} x_t + W_{hc} *_{\mathcal{G}} h_{t-1} + b_c) \\ o &= \sigma(W_{xo} *_{\mathcal{G}} x_t + W_{ho} *_{\mathcal{G}} h_{t-1} + w_{co} \odot c_{t-1} + b_o) \\ h_t &= o_t \odot \tanh(c_t) \end{aligned} \quad (4)$$

Due to the changing nature of the graph adjacency network and the limitations of training across multiple time periods, the stateless graph convolutional LSTM model (S-GC-LSTM) is harnessed, as shown in Equation 5.

$$\begin{aligned}
i &= \sigma(W_{xi} *_{\mathcal{G}} x_t) \\
f &= \sigma(W_{xf} *_{\mathcal{G}} x_t) \\
c_t &= i_t \odot \tanh(W_{xc} *_{\mathcal{G}} x_t + b_c) \\
o &= \sigma(W_{xo} *_{\mathcal{G}} x_t + b_o) \\
h_t &= o_t \odot \tanh(c_t)
\end{aligned} \tag{5}$$

This modified GC-LSTM takes in a single time point graph with an initial cell and hidden state set to 0, thus removing much of the complexity from a *typical* GC-LSTM model.

## Experimental Study: Causality using GNNs

This section presents the deployment of AINET to the trade dataset, and explores its results and efficacy.

### Data Setup

The data used in this manuscript are from the United Nations Commodity Trade Statistics Database (UN Comtrade), and consists of monthly international trade data on soybeans; specifically the harmonized system (HS) code 1201 - "soya beans; whether or not broken" (UN Comtrade 2021). Variables include trade value in dollars and net weight in kilograms imported or exported for bilateral trade around the world. Data are downloaded using UN Comtrade's API and processed using Pandas. First, the data are filtered to only include trade data with at least a net weight of 50 kg and a trade value of \$70,000. Due to limitations of trade data availability for important countries, export data for soybean trade from country A to country B is denoted as import trade to country B from country A. This ensures that country B's lack of reporting does not significantly affect the structure of the graph. In cases that both converted export data and import data were available for a single node  $n$  at time  $t$ , the average of the two is taken. The final modeling dataset consists of 70 months of data starting in October 2014 extending to July 2020 for 50 countries.

Figure 1 shows the structure of the trade network for soybeans in October 2017 (before any relevant outlier event) with a strong connection between the U.S. and China. Nodes, in red, are sized based on the net weight of soybeans imported at time  $t$ , the transparency and width of directed edges show the trade value of soybeans exported from a country such as Brazil and imported to a country like China.

Graph modeling data is then prepared with a node dictionary containing two features: the net weight at time  $t$  and the soybean trade unit value ( $TUV$ ) at time  $t + 1$ . As shown in Equation 6,  $TUV_{n,t}$  is calculated by taking the sum of all partner trade value imports ( $TVI$ ) for node  $n$  at time  $t$ . Trading partners for each node are filtered to the 10 largest trading partners by trade value.

$$TUV_{n,t} = \sum_{p=1}^P (TVI_{n,p,t}) \tag{6}$$

The trade value at time  $t$  are added to a directed-weighted adjacency list as a tuple (importer index, exporter index, trade value). This procedure was repeated for all 70 months of available data.

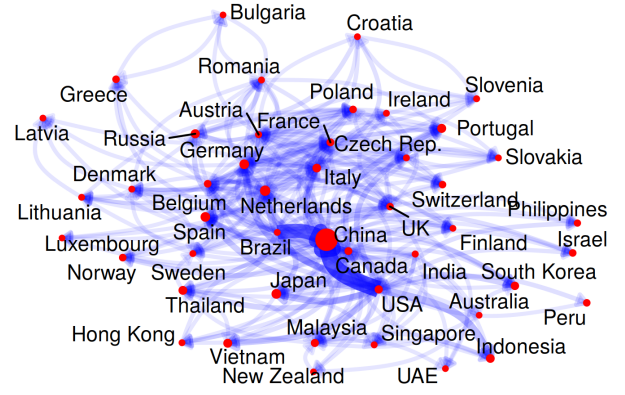


Figure 1: Soybean trade flows in Oct. 2017

### Experimental Setup

GNN modeling was conducted using the PyTorch Geometric Temporal (Benedek and Paul 2020) package. A variety of neural networks model structures were considered including GCN, GC-LSTM, and the S-GC-LSTM. Networks with multiple layers of connected cells were considered along with single cells. Outputs from graph layer(s) are combined in a linear layer for trade prediction. The network was *optimized* using the stochastic gradient descent algorithm ADAM with adjustments for the learning rate (1, 0.1, 0.01) and weight decay (0.05, 0.01, 0.00) (Kingma and Ba 2017). Stochastic gradient descent is generally used in signal processing applications to distinguish "signal" from "noise"; a notion we deemed relevant to trade values (i.e. identifying *major* trade partners for clearer cause-effect outcomes).

A variety of hyperparameters were tuned including the filter size  $K$  (2, 3, and 5), dropout (0, 0.1, 0.25, 0.5), model layers (1, 3, 5), and activation function for hidden layers ( $\tanh$ , leaky  $ReLU$ ,  $ReLU$ , and  $\text{sigmoid}$ ). These model options were trained to 100 epochs, and the performance compared to determine the optimal set of parameters for longer training. Given the "time-series" nature of the dataset, train-test splits were analyzed at the final month, 6 months, and 24 months. The 6 month model is used going forward.

To contrast using traditional tools, we also developed forecasts using ARIMA models. Economic forecasting typically uses auto regressive models as a baseline for prediction performance (Carriero, Galvão, and Kapetanios 2019). Univariate auto regressive models highlight the relationship of previous time periods and its influence on forecasting  $t + 1$ . Since the predictive power of ARIMA is strong for forecasting commodity goods, as highlighted in (Batarseh et al. 2019), the model selected was ARIMA(0,1,1): with one first difference, and one moving average for every country. ARIMA models were trained and then forecasted on 1 month, 6 months, and 24 months forward. ARIMA models could be adjusted to accommodate the data for each country time series, however, it was decided to use the most common lag, difference, and moving average terms across models for this exercise to avoid any bias. Additional details of our data and model setup can be found in our AINET repository <sup>1</sup>.

<sup>1</sup><https://github.com/AndersonMonken/AINET-GNN-Trade21>

## Experimental Results

The modified S-GC-LSTM model was determined to be the strongest GNN. Though model tuning yielded similar results, the best hyperparameters for the S-GC-LSTM model were learning rate at 0.1 with no weight decay, dropout rate of 0.1, a  $\tanh$  activation function of hidden layers, filter size  $K$  of 3 for graph convolutions, and five layers in the model. The first layer in the model takes in the single feature input for each node and outputs a 16 channel tensor, the second layer transformation  $16 \rightarrow 32$ , the third layer  $32 \rightarrow 64$ , the fourth layer  $64 \rightarrow 32$ , and the fifth graph layer  $32 \rightarrow 16$ . The 16 channels are then inputs to a linear regression prediction. The training epoch graph in Figure 2 shows the consistent improvement in performance with increasing training time.

Figure 3 plots the  $TUV$  for China, France, the U.S., and for the global aggregate for soybean trade. There are lines for the actual  $TUV$  (red), the training prediction (blue), and the testing prediction (green). The results of the global aggregate  $TUV$  can be interpreted as the potential identification of the  $TUV$  common factor. The predicted  $TUV$  for the U.S. indicates mostly a stable trend during 2017-2019, barring a few spikes coincidental with failed attempts to reconcile U.S.-China trade relations. Common factors are typically identified with dynamic factor models, such as how (Luciani 2020) identified a less volatile measure of U.S. core inflation. The AINET framework is predicting a common factor series that captures the true overall trend of global soybean  $TUV$  over time and reduces the series volatility. Price and value of commodity goods is historically a volatile series that is difficult to accurately forecast. However, the AINET framework is able to reveal the direction of soybean  $TUV$  discounting idiosyncratic shocks.

The results from the baseline ARIMA models showed greater forecasting accuracy when compared to GNN models, as expected, however they don't represent the network of flows, address causality, account for outliers, or present support for public policy making. ARIMA and parametric Ordinary Least Squares (OLS) models are part of the initial stages of understanding trade trend-lines, however, they don't provide any support beyond that.

The relationship analysis of the economic process requires understanding of theoretical macroeconomics and trade concepts; and so to improve predictions and analyze causality for soybean  $TUV$ , a form of one of these models is used. However, the extensive parametric knowledge necessary to understand trade using traditional economic methods can make the conversation of international trade more exclusive to the field of economics. However, AINET provides insight into trade that pair up AI scientists to economists.

### Assurance of GNN

AI assurance is a process that is applied at all stages of the AI engineering lifecycle ensuring that any intelligent system is producing outcomes that are valid, verified, data-driven, trustworthy, and explainable to a layman, ethical in the context of its deployment, and unbiased in its learning. Validating and verifying AI algorithms is a critical component in any AI deployment. In this section, we discuss measures we




Figure 2: Model training loss curve on train and test datasets

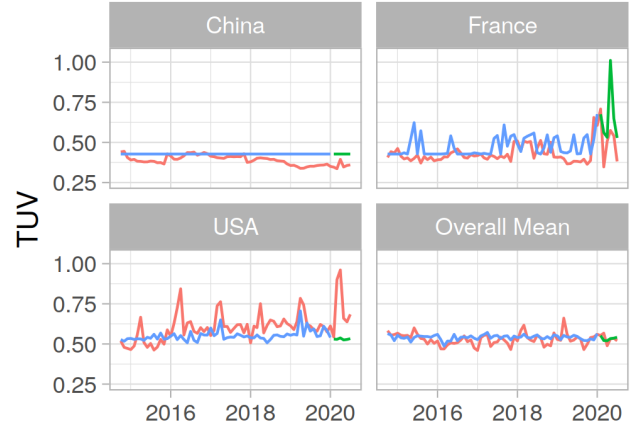


Figure 3: TUV predictions for selected countries

took to verify the algorithms.

Neural networks, including GNNs, have learning drawbacks, such as gradient vanishing and exploding. For instance, neural networks cannot process very long sequences if using certain activation functions such as the rectified linear unit. Accordingly, we opted to using LSTMs, as they are a modified version of neural networks that make it easier to remember past data in memory, for the following reasons: (1) To avoid learning bias issues with the GNN, and (2) To allow the GNN to learn causality beyond a limited time frame. For verification of learning quality and optimal selection of hyperparameters, multiple measures were collected: 1. Root Mean Square Error (RMSE), 2. Mean Absolute Percent Error (MAPE), and 3. Mean Absolute Error (MAE).

RMSE is the loss metric used to train AINET, Equation 7, where  $\hat{y}_{n,t}$  is the prediction of node  $n$  at time  $t$  and  $y_{n,t}$  is the actual value. The difference is aggregated and divided by total time periods  $T$  and total number of nodes  $N$ .

$$RMSE = \sqrt{\frac{\sum_{t=1}^T \left( \sum_{n=1}^N (\hat{y}_{n,t} - y_{n,t})^2 \right)}{NT}} \quad (7)$$

All three measures are captured in two dimensions: by year, and by country. Both aspects consistently reported lower training error rates than testing error rates, which provides statistical confidence in the predictions presented. Performing assurance measures that are relevant to trade were also crucial. For instance, cases of *irrelevant* outlier trade events were removed such as the removal of Belarus and Es-

tonia due to unusually high or low  $TUV$ . Similarly, China stopped reporting its official imports of soybeans to the national community in year 2017, which lead to estimating China’s imports through exports from other countries. Moreover, as it is stated in (Gopinath, Batarseh, and Beckman 2020), using major world traders for predictions and evaluating causality proved better than using all countries. Minor world traders of soybeans deviate the model from learning accurately, for instance, the trade of Lithuania is not as influential to causality such as Brazil or Japan. A summary of top soybean traders’ RMSE and MAPE values is shown in Table 1.

Country	Training MAPE	Test MAPE	Train RMSE	Test RMSE
U.S.	0.106	0.246	0.085	0.249
Australia	0.276	0.231	0.438	0.259
Brazil	0.648	0.850	0.220	0.257
Canada	0.277	0.315	0.126	0.162
China	0.100	0.202	0.045	0.073
France	0.147	0.342	0.084	0.212
Germany	0.072	0.111	0.034	0.044
India	0.289	0.237	0.159	0.140
Indonesia	0.196	0.395	0.125	0.155
Italy	0.089	0.156	0.042	0.058
Japan	0.075	0.041	0.053	0.024
UK	0.134	0.137	0.088	0.057

Table 1: Model Performance from AINET.

In the next section, we explore GNNs usage for public policy before/during/after two major outlier events: the U.S.-China trade war as well as Covid-19.

### Policy Relevance of AINET

Efficient and nimble agriculture and food industries are vital to human survival. In the past three years, global agriculture has been buffeted by many shocks. Such unprecedented uncertainties have affected the range of decisions starting at the farm and culminating at the consuming household or ports (exports/imports). Decisions to plant, maintain crop progress, harvest, and market in the near-term and to invest in farm assets in the medium-term have all been impacted by serious supply-side disruptions (e.g. flooding or drought), significant uncertainty in demand (e.g. U.S.-China, China-Australia trade wars) and sudden collapse of both supply of inputs and demand for output (e.g. Covid-19). Considering that the global economy was in an equilibrium with trade flows minimizing any arbitrage opportunities, these disruptions have the potential to either shut down an edge or make a node smaller than before. The GNNs offer the opportunity to visualize, measure, and understand the alternative routes for trade in search of new opportunities. Of particular interest would be the trade diversion effects (e.g. transiting products through Vietnam during U.S.-China trade war) as well as the price ( $TUV$ ) adjustment to reach a new equilibrium.

#### U.S. - China Trade War 2018

Soybean farmers in 2018 were experiencing tremendous challenges due to the escalating trade war between the U.S. and China. In 2017, 57% of all U.S. soybean exports went

to China for a value of \$12.2 billion dollars (UN Comtrade 2021). The U.S. soybean industry suffered greatly in 2018, with exports to China declining to \$3.1 billion dollars, but the damage could have been far greater if a trade deal had not been reached in 2019. As shown in Figure 4b for October 2018, there is significantly lower trade between the U.S. and China (highlighted in red where we would expect that edge) compared to Figure 4a from a year prior. Using AINET, we consider the following question: *What would the effect have been if soybean trade had been cut off entirely between the U.S. and China in 2019?*

Using AINET, we drop all soybean trade between the U.S. and China for the year 2019 to demonstrate the counterfactual capabilities of the S-GC-LSTM method. The  $TUV$  is modeled for all nodes to determine the effect to the U.S. and the rest of the global soybean industry. We find from AINET that the impact to the U.S. if trade in soybeans had ended in China for 2019 would have resulted in a stable  $TUV$  for China, a spike and return to the long-term trend of  $TUV$  for Brazil, lower  $TUV$  for the U.S., and mixed effects for the rest of the world. The stability of  $TUV$  for China along with falling  $TUV$  for U.S. likely created arbitrage opportunities for countries such as Canada (shown in Figure 4a) and Vietnam. Note the Chinese edge originating from Canada significantly thickened in Figure 4a relative to that in Figure 1. While Vietnam was expected to transship soybean to China with its actual imports doubling during the trade war months, Canada’s near tripling of soybean imports from U.S. was less noticed. The Brazilian spike in  $TUV$  and return towards the long-term trend are likely due to Yuan’s appreciation relative to the Brazilian Real during the trade war months. The use of AINET uncovered that U.S. soybean exports have diversified significantly due to the trade war, which is an advantage in terms of economic independence and stability for the American farmer.

#### Trade During Covid-19

The pandemic has had varying effects on trade flows of soybeans. Figure 4c shows that the spread of Covid-19 had thinned the edges especially with U.S. and European countries, which were shut down in early 2020. However, with China recovering faster from Covid-19 and Brazil remaining open despite the pandemic, the former returned to importing a significant level of Brazilian soybeans. The rapid devaluation of the Brazilian Real and lack of national lockdowns across Brazil during this time period likely influenced Brazil becoming the preferred trading partner in the soybean sector, as analyzed in (Ustinova 2020). Additionally, Brazil’s soybean production is usually challenged with seasonal harvest rain, however, in 2020, that *usual* challenge was not evident. Meanwhile the U.S. was in the midst of lockdowns that impacted commodity supply chains in April which resulted reduced soybean trade, and eventually harmed American soybean farmers and producers.

Figure 4c depicts the interaction described between China and Brazil. The red dot for China shows the spike in soybean imports during April compared to imports in 2017 shown in Figure 1. Likewise, India’s  $TUV$  fell as a result of its lower tariffs during the COVID-19 surge, coinciding with a seven-






Figure 4: Soybean trade flows for major countries in soybeans global trade

fold increase in soybean imports relative to the period before the pandemic. In this case, the increased trade resulted in an overall considerable increase in soybean trade value between China and Brazil as well as India and US. AINET recognizes this causal phenomenon with a decrease in  $TUV$  for overall soybean trade (Figure 3).

## Conclusion

In this paper we presented the AINET method to study international trade flows for causal analysis. AINET is based on GNNs; however, we have modified multiple parts of the conventional GNN process to capture data shifts, provide AI explainability through network graphs, increase algorithmic accuracy, reduce bias, and measure the effects of international outlier events on data. Our learned graph model can capture fast-paced changes in global markets to aid public policy making and eventually lead to lower disruptions to international trade flows.

## Acknowledgements

We thank Danilo Cascardi-Garcia and Theresa Dinh for their helpful feedback and comments. We also thank Jennifer Massey for her professional guidance, feedback, and ongoing support for this research.

## References

- Anderson, J. E. 1979. A theoretical foundation for the gravity equation. *The American economic review* 69(1):106–116.
- Batarseh, F., and Yang, R. 2017. *Federal Data Science: Transforming Government and Agricultural Policy using Artificial Intelligence*. Elsevier’s Academic Press.
- Batarseh, F.; Gopinath, M.; Nalluru, G.; and Beckman, J. 2019. Application of machine learning in forecasting international trade trends.
- Benedek, R., and Paul, S. 2020. PyTorch Geometric Temporal.
- Bhattacharya, K.; Mukherjee, G.; and Manna, S. S. 2007. The international trade network. *New Economic Windows* 139–147.
- Carriero, A.; Galvão, A. B.; and Kapetanios, G. 2019. A comprehensive evaluation of macroeconomic forecasting methods. *International Journal of Forecasting* 35(4):1226–1239.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2017. Convolutional neural networks on graphs with fast localized spectral filtering.
- Elman, J. L. 1990. Finding structure in time. *Cognitive science* 14(2):179–211.
- Ghoddusi, H.; Creamer, G. G.; and Rafizadeh, N. 2019. Machine learning in energy economics and finance: A review. *Energy Economics* 81(C):709–727.
- Gopinath, M.; Batarseh, F. A.; and Beckman, J. 2020. Machine learning in gravity models: An application to agricultural trade. Working Paper 27151, National Bureau of Economic Research.
- Gopinath, M.; Batarseh, F. A.; Beckman, J.; Kulkarni, A.; and Jeong, S. 2021. International agricultural trade forecasting using machine learning. *Data amp; Policy* 3:e1.
- Graves, A. 2014. Generating sequences with recurrent neural networks.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kim, Y. 2014. Convolutional neural networks for sentence classification.
- Kingma, D. P., and Ba, J. 2017. Adam: A method for stochastic optimization.
- Lecun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86:2278 – 2324.
- Luciani, M. 2020. Common and idiosyncratic inflation. *FEDS Note. Board of Governors of the Federal Reserve System*.
- Mullainathan, S., and Spiess, J. 2017. Machine learning: An applied econometric approach. *Journal of Economic Perspectives* 31(2):87–106.
- Panford-Quainoo, K.; Bose, J.; and Defferrard, M. 2020. Bilateral trade modeling with graph neural networks.
- S. 1558 - AIIA. 2020. Artificial Intelligence Initiative Act. *U.S. Congressional Senate Committee on Commerce, Science, and Transportation*.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Kin Wong, W.; and Chun Woo, W. 2015. Convolutional lstm network: A machine learning approach for precipitation nowcasting.
- UN Comtrade. 2021. United nations department of economic and social affairs / statistics division. united nations commodity trade statistics database.
- Ustinova, W. 2020. Brazilian agricultural sector thrives despite covid-19 pandemic. *GAIN report. U.S. Department of Agriculture, Foreign Service Agriculture*.
- Wohl, I., and Kennedy, J. 2018. Neural network analysis of international trade.
- Yamaguchi, K.; Sakamoto, K.; Akabane, T.; and Fujimoto, Y. 1990. A neural network for speaker-independent isolated word recognition. In *First International Conference on Spoken Language Processing*.